

Position Paper: Goals of the Luau Type System

LILY BROWN, ANDY FRIESEN, and ALAN JEFFREY, Roblox, USA

Luau is the scripting language that powers user-generated experiences on the Roblox platform. It is a statically-typed language, based on the dynamically-typed Lua language, with type inference. These types are used for providing editor assistance in Roblox Studio, the IDE for authoring Roblox experiences. Due to Roblox’s uniquely heterogeneous developer community, Luau must operate in a somewhat different fashion than a traditional statically-typed language. In this paper, we describe some of the goals of the Luau type system, focusing on where the goals differ from those of other type systems.

ACM Reference Format:

Lily Brown, Andy Friesen, and Alan Jeffrey. 2021. Position Paper: Goals of the Luau Type System. In *HATRA ’21: Human Aspects of Types and Reasoning Assistants*. ACM, New York, NY, USA, 7 pages.

1 INTRODUCTION

The Roblox platform allows anyone to create shared, immersive, 3D experiences. As of July 2021, there are approximately 20 million experiences available on Roblox, created by 8 million developers [19]. Roblox creators are often young: there are over 200 Roblox kids’ coding camps in 65 countries listed by the company as education resources [18]. The Luau programming language [17] is the scripting language used by creators of Roblox experiences. Luau is derived from the Lua programming language [7], with additional capabilities, including a type inference engine.

This paper will discuss some of the goals of the Luau type system, such as supporting goal-driven learning, non-strict typing semantics, and mixing strict and non-strict types. Particular focus is placed on how these goals differ from traditional type systems’ goals.

2 NEEDS OF THE ROBLOX PLATFORM

2.1 Heterogeneous developer community

Need: *a language that is powerful enough to support professional users, yet accessible to beginners*
Quoting a Roblox 2020 report [16]:

- *Adopt Me!* now has over 10 billion plays and surpassed 1.6 million concurrent users earlier this year.
- *Piggy*, launched in January 2020, has close to 5 billion visits in just over six months.
- There are now 345,000 developers on the platform who are monetizing their games.

This demonstrates the heterogeneity of the Roblox developer community: developers of experiences with billions of plays are on the same platform as children first learning to code. Both of these groups are important to support: the professional development studios bring high-quality experiences to the platform, and the beginning creators contribute to the energetic creative community, forming the next generation of developers.

2.2 Goal-driven learning

Need: *organic learning for achieving specific goals*

All developers are goal-driven, but this is especially true for learners. A learner will download Roblox Studio (the creation environment for the Roblox platform) with an experience in mind, such as designing an obstacle course to play in with their friends.



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

HATRA ’21, October 2021, Chicago, IL

© 2021 Roblox.

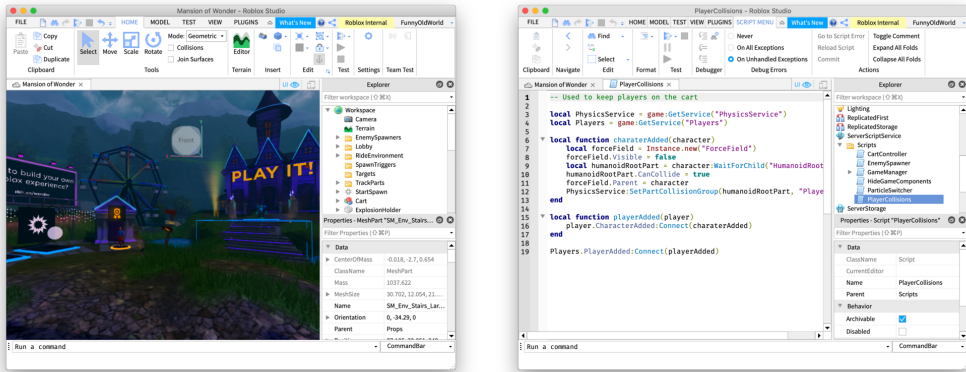


Fig. 1. Roblox Studio’s 3D environment editor (a), and script editor (b)

The user experience of developing a Roblox experience is primarily a 3D interactive one, seen in Fig. 1(a). The user designs and deploys 3D assets such as terrain, parts and joints, providing them with physics attributes such as mass and orientation. The user can interact with the experience in Studio, and deploy it to a Roblox server so anyone with the Roblox app can play it. Physics, rendering and multiplayer are all immediately accessible to creators.

At some point during experience design, the experience creator has a need that can’t be met by the game engine alone, such as “the stairs should light up when a player walks on them” or “a firework is set off every few seconds”. At this point, they will discover the script editor, seen in Fig. 1(b).

This onboarding experience is different from many initial exposures to programming, in that by the time the user first opens the script editor, they have already built much of their creation, and have a very specific concrete aim. As such, Luau must allow users to perform a specific task with as much help as possible from tools.

A common workflow for getting started is to Google the task, then cut-and-paste the resulting code, adapting it as needed. Since this is so common, backward compatibility of Luau with existing code is important, even for learners who do not have an existing code base to maintain.

Type-driven tools are useful to all creators, in as much as they help them achieve their current goals. For example type-driven autocomplete, or type-driven API documentation, are of immediate benefit. Traditional typechecking can be useful, for example for catching spelling mistakes, but for most goal-driven developers, the type system should help or get out of the way.

2.3 Type-driven development

Need: *a language that supports large-scale codebases and defect detection*

Professional development studios are also goal-directed (though the goals may be more abstract, such as “decrease user churn” or “improve frame rate”) but have additional needs:

- *Code planning*: code spends much of its time in an incomplete state, with holes that will be filled in later.
- *Code refactoring*: code evolves over time, and it is easy for changes to break previously-held invariants.
- *Defect detection*: code has errors, and detecting these at runtime (for example by crash telemetry) can be expensive and recovery can be time-consuming.

Detecting defects ahead-of-time is a traditional goal of type systems, resulting in an array of techniques for establishing safety results, surveyed for example in [13]. Supporting code planning and refactoring are some of the goals of *type-driven development* [1] under the slogan “type, define, refine”. A common use of type-driven development is renaming a property, which is achieved by changing the name in one place, and then fixing the resulting type errors—once the type system stops reporting errors, the refactoring is complete.

To help support the transition from novice to experienced developer, types are introduced gradually, through API documentation and type discovery. Type inference provides many of the benefits of type-driven development even to creators who are not explicitly providing types.

3 GOALS OF THE TYPE SYSTEM

3.1 Infallible types

Goal: *provide type information even for ill-typed or syntactically invalid programs.*

Programs spend much of their time under development in an ill-typed or incomplete state, even if the final artifact is well-typed. If tools such as autocomplete and API documentation are type-driven, this means that tooling needs to rely on type information even for ill-typed or syntactically invalid programs. An analogy is infallible parsers, which perform error recovery and provide an AST for all input texts, even if they don’t adhere to the parser’s syntax.

Program analysis can still flag type errors, which may be presented to the user with red squiggly underlining. Formalizing this, rather than a judgment $\Gamma \vdash M : T$, for an input term M , there is a judgment $\Gamma \vdash M \Rightarrow N : T$ where N is an output term where some subterms are *flagged* as having type errors, written \underline{N} . Write $\text{erase}(N)$ for the result of erasing flaggings: $\text{erase}(\underline{N}) = \text{erase}(N)$.

For example, in Lua, the `string.find` function expects two strings, and returns the offsets for that string:

$$\text{string.find}(\text{“hello”}, \text{“ell”}) \rightarrow (2, 4) \quad \text{string.find}(\text{“world”}, \text{“ell”}) \rightarrow (\text{nil}, \text{nil})$$

and in Luau it has the type:

$$\text{string.find} : (\text{string}, \text{string}) \rightarrow (\text{number?}, \text{number?})$$

In a conventional type system, there is no judgment for ill-typed terms such as `string.find(“hello”, 37)` but in an infallible system we flag the error and approximate the type, for example:

$$\vdash \text{string.find}(\text{“hello”}, 37) \Rightarrow \underline{\text{string.find}(\text{“hello”}, 37)} : (\text{number?}, \text{number?})$$

The goal of infallible types is that every term has a typing judgment given by flagging ill-typed subterms:

- *Typability*: for every M and Γ , there are N and T such that $\Gamma \vdash M \Rightarrow N : T$.
- *Erasure*: if $\Gamma \vdash M \Rightarrow N : T$ then $\text{erase}(M) = \text{erase}(N)$

Some issues raised by infallible types:

- Which heuristics should be used to provide types for flagged programs? For example, could one use minimal edit distance to correct for spelling mistakes in field names?
- How can we avoid cascading type errors, where a developer is faced with type errors that are artifacts of the heuristics, rather than genuine errors?
- How can the goals of an infallible type system be formalized?

Related work: there is a large body of work on type error reporting (see, for example, the survey in [4, Ch. 3]) and on type-directed program repair (see, for example, the survey in [8, Ch. 3]), but less on type repair. The closest work is Hazel’s [11] *typed holes* where \underline{N} is treated as a partially-filled hole in the program, though in that work partially-filled holes are not erased at run-time. Many

compilers perform error recovery during typechecking, but do not provide a semantics for programs with type errors.

3.2 Strict types

Goal: *no false negatives*.

For developers who are interested in defect detection, Luau provides a *strict mode*, which acts much like a traditional, sound, type system. This has the goal of “no false negatives” where any possible run-time error is flagged. This is formalized using:

- *Operational semantics*: a reduction judgment $M \rightarrow N$ on terms.
- *Values*: a subset of terms representing a successfully completed evaluation.

Error states at runtime are represented as stuck states (terms that are not values but cannot reduce), and showing that no well-typed program is stuck. This is not true if typing is infallible, but can fairly straightforwardly be adapted. We extend the operational semantics to flagged terms, where $M \rightarrow M'$ implies $\underline{M} \rightarrow \underline{M}'$, and for any value V we have $\underline{V} \rightarrow V$, then show:

- *Progress*: if $\vdash M \Rightarrow N : T$, then either $N \rightarrow N'$ or N is a value or N has a flagged subterm.
- *Preservation*: if $\vdash M \Rightarrow N : T$ and $N \rightarrow N'$ then $M \rightarrow^* M'$ and $\vdash M' \Rightarrow N' : T$.

For example in typechecking the program:

$$\text{local } (i, j) = \text{string.find}(x, y); \text{ if } i \text{ then print}(j - i) \text{ end}$$

the interesting case is $i - j$ in a context where i has type number (since it is guarded by the if) but j has type number?. Since subtraction has type $(\text{number}, \text{number}) \rightarrow \text{number}$, this is a type error, so the relevant typing judgment is:

$$x : \text{string}, y : \text{string} \vdash (\text{local } (i, j) = \text{string.find}(x, y); \text{ if } i \text{ then print}(j - i) \text{ end}) \\ \Rightarrow (\text{local } (i, j) = \text{string.find}(x, y); \text{ if } i \text{ then print}(\underline{j - i}) \text{ end})$$

Some issues raised by soundness for infallible types:

- How should the judgments and their metatheory be set up?
- How should type inference and generic functions be handled?
- Is the operational semantics of flagged values ($\underline{V} \rightarrow V$) the right one?

Related work: gradual typing and blame analysis, e.g. [2, 21, 23]. The main difference between this approach and that of migratory typing [22] is that (due to backward compatibility with existing Lua) we cannot introduce extra code during migration.

3.3 Nonstrict types

Goal: *no false positives*.

For developers who are not interested in defect detection, type-driven tools and techniques such as autocomplete, API documentation and type-driven refactoring are still useful. For such developers, Luau provides a *nonstrict mode*, which we hope will eventually be useful for all developers. This non-strict typing mode is particularly useful when adopting Luau types in pre-existing code that was not authored with the type system in mind. Non-strict mode does *not* aim for soundness, but instead has the goal of “no false positives”, in the sense that any flagged code is guaranteed to produce a runtime error when executed.

Our previous example was, in fact, a false positive since a programmer can make use of the fact that $\text{string.find}(x, y)$ is either nil in both results or neither, so if i is non-nil then so is j . This is discussed in the English-language documentation but not reflected in the type. So flagging $(i - j)$ is a false positive.

On the face of it, detecting all errors without false positives is undecidable, since a program such as (if $f()$ then error end) will produce a runtime error when $f()$ is true. Instead we can aim for a weaker property: that all flagged code is either dead code or will produce an error. Either of these is a defect, so deserves flagging, even if the tool does not know which reason applies.

We can formalize this by defining an *evaluation context* $\mathcal{E}[\bullet]$, and saying M is *incorrectly flagged* if it is of the form $\mathcal{E}[V]$. We can then define:

- *Correct flagging*: if $\vdash M \Rightarrow N : T$ then N is correctly flagged.

Some issues raised by nonstrict types:

- Will nonstrict types result in errors being flagged in function call sites rather than definitions?
- In Luau, ill-typed property update of most tables succeeds (the property is inserted if it did not exist), and so functions which update properties cannot be flagged. Can we still provide meaningful error messages in such cases?
- Does nonstrict typing require whole program analysis, to find all the possible types a property might be updated with?
- The natural formulation of function types in a nonstrict setting is that of [6]: if $f : T \rightarrow U$ and $f(V) \rightarrow^* W$ then $V : T$ and $W : U$. This formulation is *covariant* in T , not *contravariant*; what impact does this have?

Related work: success types [6] and incorrectness logic [10].

3.4 Mixing types

Goal: *support mixed strict/nonstrict development*.

Like every active software community, Roblox developers share code with one another constantly. First- and third-party developers alike frequently share entire software packages written in Luau. To add to this, many Roblox experiences are authored by a team. It is therefore crucial that we offer first-class support for mixing code written in strict and nonstrict modes.

Some questions raised by mixed-mode types:

- How much feedback can we offer for a nonstrict script that is importing strict-mode code?
- In strict mode, how do we talk about values and types that are drawn from nonstrict code?
- How can we combine the goals of strict and nonstrict types?
- Can we have strict and non-strict mode infer the same types, only with different flagging?
- Is strict-mode code sound when it relies on non-strict code, which has weaker invariants?
- How can we avoid introducing function wrappers in higher-order code at the strict/nonstrict boundary?

Related work: there has been work on interoperability between different type systems, notably [12], but there the overall goals of the systems were similar safety properties. In our case, the two type systems have different goals.

3.5 Type inference

Goal: *infer types to allow gradual adoption of type annotations*.

Since backward compatibility with existing code is important, we have to provide types for code without explicit annotations. Moreover, we want to make use of type-directed tools such as autocomplete, so we cannot adopt the common strategy of treating all untyped variables as having type any. This leads us to type inference.

To make use of type-driven technologies for programs without explicit type annotations, we use a type inference algorithm. Since Luau includes System F, type inference is undecidable [15], but we can still make use of heuristics such as local type inference [14].

It remains to be seen if type inference can satisfy the goals of strict and non-strict types. The current Luau system infers different types in the two modes, which is unsatisfactory as it makes changing mode a non-local breaking change. In addition, non-strict inference is currently too imprecise to support type-directed tools such as autocomplete.

Some questions raised by type inference:

- How many cases in strict mode cannot be inferred by the type inference system? Minimizing this kind of error is desirable, to make the type system as unobtrusive as possible.
- Can something like the Rust traits system [5] or Haskell classes [3] be used to provide types for overloaded operators, without hopelessly confusing learners?
- Type inference currently infers monotypes for unannotated functions, in contrast to Quick-Look [20], which can infer generic types. Will this be good enough for idiomatic Luau scripts?
- Can type inference be used to infer the same types in strict and nonstrict mode, to ease migrating between modes, with the only difference being error reporting?

Related work: there is a large body of work on type inference, largely summarized in [13].

4 CONCLUSIONS

In this paper, we have presented some of the goals of the Luau type system, and how they map to the needs of the Roblox creator community. We have also explored how these goals differ from traditional type systems, where it is necessary to accommodate the unique needs of the Roblox platform. We have sketched what a solution might look like; all that remains is to draw the owl [9].

REFERENCES

- [1] Edwin Brady. 2017. *Type-Driven Development with Idris*. Manning.
- [2] Robert B. Findler and Matthias Felleisen. 2002. Contracts for Higher-order Functions. In *Proc. Int. Conf. Functional Programming*. 48–59.
- [3] Cordelia V. Hall, Kevin Hammond, Simon L. Peyton Jones, and Philip L. Wadler. 1996. Type Classes in Haskell. *ACM Trans. Program. Lang. Syst.* 18, 2 (1996), 109–138.
- [4] Bastiaan J. Heeren. 2005. *Top Quality Type Error Messages*. Ph.D. Dissertation. U. Utrecht.
- [5] Steve Klabnik, Carol Nichols, and the Rust Community. 2021. The Rust Programming Language. <https://doc.rust-lang.org/book/>
- [6] Tobias Lindahl and Konstantinos Sagonas. 2006. Practical Type Inference Based on Success Typings. In *Proc. Int. Conf. Principles and Practice of Declarative Programming*. 167–178.
- [7] Lua.org and PUC-Rio. 2021. The Lua Programming Language. <https://lua.org>
- [8] Bruce J. McAdam. 2002. *Repairing Type Errors in Functional Programs*. Ph.D. Dissertation. U. Edinburgh.
- [9] Know Your Meme. 2010. How To Draw An Owl. <https://knowyourmeme.com/memes/how-to-draw-an-owl>
- [10] Peter W. O’Hearn. 2020. Incorrectness Logic. In *Proc. Symp. Principles of Programming Languages*. Article 10, 32 pages.
- [11] Cyrus Omar, Ian Voysey, Ravi Chugh, and Matthew Hammer. 2019. Live Functional Programming with Typed Holes. In *Proc. Symp. Principles of Programming Languages*. 14:1–14:28.
- [12] Daniel Patterson and Amal Ahmed. 2017. Linking Types for Multi-Language Software: Have Your Cake and Eat It Too. In *Proc. Summit on Advances in Programming Languages*.
- [13] Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press.
- [14] Benjamin C. Pierce and David N. Turner. 2000. Local Type Inference. *ACM Trans. Program. Lang. Syst.* 22, 1 (2000), 1–44.
- [15] Partial polymorphic type inference is undecidable. 1985. Hans-J. Boehm. In *Proc. Symp. Foundations of Computer Science*. 339–345.
- [16] Roblox. 2020. Roblox Developers Expected to Earn Over \$250 Million in 2020; Platform Now Has Over 150 Million Monthly Active Users. <https://corp.roblox.com/2020/07/roblox-developers-expected-earn-250-million-2020-platform-now-150-million-monthly-active-users/>
- [17] Roblox. 2021. The Luau Programming Language. <https://luau-lang.org>
- [18] Roblox. 2021. Roblox Education: All Educators. <https://education.roblox.com/en-us/educators>
- [19] Roblox. 2021. What is Roblox. <https://corp.roblox.com>

- [20] Alejandro Serrano, Jurriaan Hage, Simon Peyton Jones, and Dimitrios Vytiniotis. 2020. A quick look at impredicativity. In *Proc. Int. Conf. Functional Programming*.
- [21] Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Proc. Scheme and Functional Programming Workshop*. 81–92.
- [22] Sam Tobin-Hochstadt, Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Ben Greenman, Andrew M. Kent, Vincent St-Amour, T. Stephen Strickland, and Asumu Takikawa. 2017. Migratory Typing: Ten Years Later. In *Proc. Summit on Advances in Programming Languages*.
- [23] Philip Wadler and Robert B. Findler. 2009. Well-typed Programs Can't be Blamed. In *Proc. European Symp. Programming*. 1–16.