Advanced topics in programming languages

Michaelmas 2025 -

$\Gamma \vdash ? : \tau$

Program synthesis

Jeremy Yallop jeremy.yallop@cl.cam.ac.uk

The program synthesis problem

What is the synthesis problem?

The problem



Challenges

Program Synthesis (Gulwani et al, 2017):

...is the task of automatically finding a program in the underlying programming language that satisfies the user intent expressed in the form of some specification.

(emphasis mine)

That is, it's a search for a constructive proof of a quantified formula:

 $\exists f. \forall input. Specification$

When is program synthesis useful?

The problem



Challenges

Efficiency in programming

(low-level code from high-level specifications)

Effective compilation

(e.g.superoptimization)

Program repair

(updating buggy programs to fit a specification)

Deobfuscation

(restoring readability)

End-user programming

(e.g. interactive programming-by-examples)

Program transformation

(updating programs as specifications evolve)

What is a specification?

The problem

Challenges

"...the user intent expressed in the form of some specification ..."

A logical specification

A type

An existing program

$$f(x,y) \ge x \land f(x,y) \ge y \qquad \begin{cases} x : \mathbb{Z} \to y : \mathbb{Z} \to \\ \{z : \mathbb{Z} \mid z = \max(x,y)\} \end{cases}$$

$$x: \mathbb{Z} \to y: \mathbb{Z} \to z: \mathbb{Z} \to z: \mathbb{Z}$$

$$slow_max(x,y)$$

Natural language

$$f(2,4) = 4, f(5,2) = 5,...$$
 "The larger of x and y"

One approach: Syntax-Guided Synthesis (SyGuS)





$$f: \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$$

$$f(x, y) = f(y, x) \land f(x, y) \ge x$$

Challenges

Reading



Example from Search-based Program Synthesis, Alur et al (2018)

 \rightarrow f(x,y) = ITE((x \leq y),y,x)

Why is program synthesis hard?

Challenge: big search space

syntactic templates

The problem

Synthesis is often based on some form of **enumeration** of programs.

Challenges

nanenges

Reading

However, the search space is extremely large (exponential in program length).

Some form of **pruning** or **guidance** is necessary, e.g. by using

abstract interpretation

domain equations component-based construction stochastic search constraint solving precise types

grammar refinement

Challenges 2: determining correctness

The problem

How can we tell when we've found a solution?

SMT solving

Type checking

Z3

 $\Gamma \vdash e : \tau$

Challenges



Human inspection



Testing



Success in limited domains

The problem

Challenges

•••

Spreadsheet formulas



Regular expressions

a(b|c)*d

SQL

queries

Trigonometric functions



Loop-free programs

> from t select * where

Bit twiddling

x & 0xBEEF << y

Background reading: Program Synthesis

The problem

Program Synthesis

Sumit Gulwani Microsoft Research sumitg@microsoft.com

Oleksandr Polozov University of Washington polozov@cs.washington.edu

> Rishabh Singh Microsoft Research risin@microsoft.com

> > ∩**⊝**₩

Boston — Delft

"This survey is a general overview of the state-of-the-art approaches to program synthesis, its applications, and subfields. We discuss the general principles common to all modern synthesis approaches such as syntactic bias, oracleguided inductive search, and optimization techniques."

Program Synthesis.

S. Gulwani, O. Polozov and R. Singh. Foundations and Trends in Programming Languages, vol. 4, no. 1-2, pp. 1–119, 2017.

Reading

0000

Online:

 $https://microsoft.com/en-us/research/wp-content/uploads/2017/10/program_synthesis_now.pdf$

Paper 1: types and examples (2015)

The problem

Challenges

Reading



Type-and-Example-Directed Program Synthesis

Peter-Michael Osera Steve Zdancewic University of Pennsylvania, USA (posera, stevez)@cis.upenn.edu



Abstract

This paper presents an algorithm for synthesizing recursive functions that process algorithm datasystes. It is founded on proof-theoretic techniques that exploit both type information and input-cupits techniques that exploit both type information and input-cupits techniques that the exploit properties that the exploit properties to extract the successful properties to extract the successful properties to expressive the substantial to the protection to make the algorithm by using a protection processing the exploit properties the exploit more than the backman to the process the exploit th

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory—Semantics; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic— Proof Theory; L.2.2 [Artificial Intelligence]: Automatic Programming—Program Synthesis

General Terms Languages, Theory

Keywords Functional Programming, Proof Search, Program Synthesis, Type Theory

1. Synthesis of Functional Programs

This spare presents a novel technique for synthesizing purely functional, recursive programs that prices algebraic datasyses. Our approach refines the venerable idea (Green 1999; Manna aprof search; rather than using just type information, our algorithm also reformation to the property of the property of the property of the search; rather than using just type information, our algorithm also in the property of the property of the property of the property of search; rather than 1990; We exploit in sext information to create a data structure, called a refinement tree, that enables efficient synthesis of non-trivial programs.

Figure 1 shows a small example of this synthesis procedure in action. For concreteness, our prototype implementation uses OCannl syntax, but the technique is not specific to that choice. The inputs to the algorithm include a type signature, the definitions of any needed auxiliary functions, and a synthesis good. In the figure, we define not

(* Type signature for natural numbers and lists *)

type nat = type list = | Nil

| S of nat | Cons of nat * list (* Good type refused by impost/cooped examples *) let stutter | list -> list |> (| | -> | | | | | | | -> | | |

(* Output: symbosized implementation of statter *)

let statter : list -> list =

let rec fi (li:list) : list =

match li with

| Nil -> 11 | Cons(n1, 12) -> Cons(n1, Cons(n1, f1 12))

Figure 1. An example program synthesis problem (above) and the resulting synthesized implementation (below).

(natural number) and 11st types without any auxiliary functions. The goal is a function named attuter of type 11st \rightarrow 11st, partially specified by a series of input-output examples given after the \mid > marker, evocative of a refinement of the goal type. The examples uggest that attuter is should produce a list that duplicates each element of the input list. The third example, for instance, means that switzer [1:0] should wish [1:1:0] be that which [1:0] should wish [1:1:0] be the standard of the

The bottom half of Figure 1 shows the output of our synthesis algorithm which is computed in negligible time (about 0.001s). Here, we see that the result is the "obvious" function that creates two copies of each Cons cell in the input list, stuttering the tail recursively via the call to fe 1 12.

General program synthesis techniques of this kind have many potential applications. Recent success stories utilize synthesis in many scenarios; programming spreadsheet macros by example (Gul-wani 2011); code completion in the context of large APIs or libraries (Perelman et al. 2012; Gwero et al. 2013); and generating cache-coherence protocols (Udupa et al. 2013), among others. In

"It is founded on proof-theoretic techniques that exploit both type information and input—output examples to prune the search space.

[...]

The goal is [...] partially specified by a series of input–output examples given after the |> marker, evocative of a refinement of the goal type.

[...]

Our algorithm [...] uses the proof-theoretic idea of searching only for programs in $\beta\text{-}$ normal, $\eta\text{-}\text{long}$ form

Paper 2: refinement types (2016)

The problem

are particularly suitable for program synthesis for two reasons. First, they offer a unique combination of expressive power and decidability, which enables automatic verification—and hence synthesis—of nontrivial programs. Second, a type-heast specification for a program can often be effectively decomposed into independent specifications for its composense, causing the synthesize to consider fewer composent confinations and leading control of the composent confinations and leading confinations are control of the composent confinations and leading control of the composent confinations and leading confination control of the composent composent confinations and leading control of the composent composent control of the composent control of

Abstract

We have evaluated our prototype implementation on a large set of synthesis problems and found that it exceeds the state of the art in terms of both scalability and usability. The tool was able to synthesize more complex programs than those reported in your work (several sorting algorithms and operations on balanced search trees), as well as most of the benchmarks tackled by existing synthesizers, often starting from a more concise and intuition, ourse insultant.

type checking, which supports specification decomposition.

We present a method for synthesizing recursive functions that

provably satisfy a given specification in the form of a poly-

morphic refinement type. We observe that such specifications

Categories and Subject Descriptors F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; 1.2.2 [Automatic Programming]: Program Synthesis

Program Synthesis from Polymorphic Refinement Types

Nadia Polikarpova Ivan Kuraj Armando Solar-Lezama
MIT CSAIL, USA
{polikarn.ivanko.asolar}@csail.mit.edu

1 Introduction

The key to scalable program synthesis is modular verification. Modularity enables the synthesize to prune candidates for different subprograms independently, whereby combinatorially reducing the size of the search space it has to consider. This explains the recent success of rpv-directed approaches to synthesis of functional programs [12, 14, 15, 22] not only do ill-typed programs wastly outnimber well-typed tones, but more former to the control of the state of the sta

Simple, coarse-grained types alone are, however, rarely sufficient to precisely describe a synthesis goal. Therefore, existing approaches supplement type information with other kinds of specifications, such as input-output examples [1, 12, 27], or pre- and post-conditions [20, 21]. Alas, the corresponding verification procedures rarely enjoy the same level of modularity as type checking, thus fundamentally limiting the scalability of

In this work we present a novel system that pushes the idea of type-directed synthesis one step further by taking adventured by of type-directed policy. It is a superior of the property of the property of refinement type [13, 33]; types decorated with predicates from a decidable logic. For example, imagine that sue resident to to synthesize the function replicate, which, given a natural transfer and as value x-produces a list that contains no summer and as value x-produces a list that contains no fixed the property of the property of the property of the property of the following signature:

 $\texttt{replicate} :: n : \texttt{Nat} \to \! x : \alpha \! \to \! \{\nu \colon \texttt{List} \; \alpha \, | \, \texttt{len} \; \nu \! = \! n \}$

"We present a method for synthesizing recursive functions that provably satisfy a given specification in the form of a polymorphic refinement type.

"a unique combination of expressive power and decidability [...] a type-based specification for a program can often be effectively decomposed into independent specifications for its components [...] leading to a combinatorial reduction in the size of the search space.

"The tool was able to synthesize more complex programs than those reported in prior work (several sorting algorithms and operations on balanced search trees) [...] often starting from a more concise and intuitive user input."



Paper 3: paramorphisms (2024)

The problem

Challenges

Reading

••••

Recursive Program Synthesis using Paramorphisms

QIANTAN HONG, Stanford University, USA

ALFX AIKFN Stanford University USA

We show that yuthesizing recursive functional programs using a class of primitive recursive combinators to shoot simpler and solves more benchmark from the literature than previously proposed approaches. Our such of synthesizes person-politions, a class of programs that includes the most common recursive programming patterns on algebras data byces. The crue of our approach is to uptile they synthesize problem into the post are multi-bole resplict that fixes the recursive structure, and a search for non-recursive program fragments to fill the remulate holes.

CCS Concepts: • Software and its engineering → General programming languages, Programming by example, Search-based software engineering; Automatic programming.

Additional Key Words and Phrases: Program Synthesis, Examples, Stochastic Synthesis, Recursion Schemes

Qiantan Hong and Alex Aiken. 2024. Recursive Program Synthesis using Paramorphisms. Proc. ACM Program. Lang. 8, PLDI, Article 151 (June 2024), 24 pages. https://doi.org/10.1145/3656381

1 INTRODUCTION

We consider the problem of synthesizing recursive programs from input-output examples. Following previous work, we consider functional programs over algebraic data types such as the natural numbers, lists, and trees [Kneuss et al. 2013; Lubin et al. 2020; Osera and Zdancewic 2015]. For example, consider a program that appends two lists:

```
append Nil l = l
append (Cons h(t) l = Cons h (append t(l))
```

This program uses general recursion, that is, the function append is explicitly recursively defined, with calls to append within its definition. Depending on what other language features are present, unrestricted general recursion is difficult to reason about; for example, proving termination of seneral recursive programs is normally non-trivial.

In practice many iterative/recursive programs, including append, can be expressed using more restricted primitive recursive constructs. The essence of primitive recursion is that the number of iterations or recursive invocations is known when the function is first called. For example, the fold combinator captures a typical primitive recursive pattern where the number of recursive calls is the length of the list argument. A standard (general recursive) definition of fold is:

"Our method synthesizes *paramorphisms*, a class of programs that includes the most common recursive programming patterns on algebraic data types.

[...]

The paramorphism combinator on lists is:

para Nil
$$g_{Nil} g_{Cons} = g_{Nil}$$

$$\mathsf{para}\ (\mathsf{Cons}\ h\ t)\ \mathsf{g}_{\mathsf{Nil}}\ \mathsf{g}_{\mathsf{Cons}} = \mathsf{g}_{\mathsf{Cons}}\ h\ (t,\ \mathsf{para}\ t\ \mathsf{g}_{\mathsf{Nil}}\ \mathsf{g}_{\mathsf{Cons}})$$

[...]

We have shown by experiment that an implementation of our approach is able to synthesize all the problems handled by the current state of the art as well as substantially harder problems."

Writing suggestions

The problem

Decidability

How does the system determine when a solution is valid?

Scalability

How complex can specifications be?

How large can generated programs be?
What subset of the language is targeted?

How long does synthesis take?

Challenges

Practicability

How easy is it for users to express specifications?

Applicability

What range of problems might the system apply to?

