Advanced topics in programming languages

Michaelmas 2024

Program synthesis $\Gamma \vdash ?: \tau$

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

The program synthesis problem

The problem

Challenges

Reading

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

 $\bigcirc \bigcirc \bigcirc$

Efficiency in programming

(low-level code from high-level specifications)

Effective compilation

(e.g.superoptimization)

Program repair

(updating buggy programs to fit a specification)

Challenges

Reading

Deobfuscation End-use

End-user programming

(restoring readability)

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

Program transformation

(updating programs as specifications evolve)

What is a specification?



One approach: Syntax-Guided Synthesis (SyGuS)



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

Why is program synthesis hard?

Challenge: big search space

The problem

Challenges

Reading

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

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 grammar refinement syntactic templates domain equations component-based construction stochastic search constraint solving precise types

Challenges 2: determining correctness



Success in limited domains





Background reading: Program Synthesis



"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."

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

Online:

Reading

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

Paper 1: types and examples (2015)



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 (assumpset. It is founded on proof-theoretic techniques that exploit both type information and input-ouput complex to pruse the succh base. The algorithm uses orfgeneous trees, a data structure that succirably represents constraints on the shape of governated code. We evaluate the algorithm use singleneous and several hone within larger examples. Our results domonstrate that the approxhem tens to competions that uses of due and the situation of the structure of the succirable size of the generated programs.

Categories and Subject Descriptors D.3.1 [Programming Langauges]; Formal Definitions and Theory—Semantics; F.4.1 [Mathconstical 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 paper presents a novel technique for synthesizing purely functional, recenvire programs that process algebraic datatypes. Our approach refines the venerable idea (Green 1999; Manna and Wahnger 1980); Otteniag programs purphesis as a kiloni of proof search: rather than using just type information, our algorithm also is of the search space. We exploit this texin 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: OCaml 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 goal. In the figure, we define nat (* Output: symbolical implementation of statute *) let statuter : list → list = let read filialist) : list = match li with | Nil→J l] | Cons(n1, 12) → Cons(n1, Cons(ni, fi 12)) in fl

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

(natural number) and 11:st types without any auxiliary functions. The goal is a function named statter of type 11:st, partially specified by a series of input-output examples given after the |> marker, evocative of a refinement of the goal type. The examples suggest that statters aboudd produce a list that duplicates each element of the input list. The third example, for instance, means that statters ($|z_10\rangle$ should yield ($|z_1i_2i_2i_1\rangle$)

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 "lobvios" function that creates two copies of each Cons cell in the input list, stuttering the tail recursively via the call to f: 1.2.

General program synthesis techniques of this kind have many potential applications. Recent success stories utilize synthesis in many scenarios: programming spreadsheet macros by example (Galwmi 2011); code completion in the context of large APIs or libraries (Perelman et al. 2012; Gvero et al. 2013); and generating cache-coherence protocols (Udupa et al. 2013), anong 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 β -normal, η -long form

Challenges

Paper 2: refinement types (2016)

Program Synthesis from Polymorphic Refinement Types

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

Abstract

We present a method for synthesizing recursive functions that provably satisfy a given specification in the form of a polymorphic refinement type. We observe that such specifications are particularly utiliable for program synthesis for two reasons. First, they offer a unique combination of expressive power and decidability, which enables automatic event decidability which enables automatic event independent specifications for its composents, cansing the synindependent specifications for its composents, cansing the synthetic combination functions in this size of the search space. At this core of our synthesis procedure is a new algorithm for refinement type checking. which surports specification decomposition

We have evaluated our prototype implementation on a large set of synthesis problems and found hat it exceeds the state of the art in terms of both scalability and usability. The tod was able to synthesize more complex programs than those reported in prior 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 innitive user input.

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

1. Introduction

The key to scalable program synthesis is modular verification. Modularity enables the synthesizer to prune candidates for different subprograms independently, whereby combinatorally reducing the size of the search space it has to consider. This explain the next success of yp-actic algoroaches to synthesis of functional programs [12, 14, 15, 27]; not only do iii) typed programs way domunibre well apped ones, but more importantly, a type error can be detected long before the whole program is part together.

Simple, coarse-grained types alone are, however, rarely sufficient to precisely describe a synthesis goal. Therefore, eristing 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 proceedure narely enjoy the same level of modularity as type hecking, thus fundamentally limiting the scalability of these techniques.

In this work, we present a novel system that pushes the idea of type-directed synthesis one step further by taking advantage of reforment type [13, 33]; types decorated with predicates from a decidable [06]; For example, imagine that a user intends to synthesize the function replicate, which, given a natural number n and a value, z produces a list that contains n copies of z. In our system, the user can express this intent by providing the following signature:

replicate :: $n: Nat \rightarrow x: \alpha \rightarrow \{\nu: List \alpha | len \nu = n\}$

"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."

Challenges

The problem

The problem

Recursive Program Synthesis using Paramorphisms

QIANTAN HONG, Stanford University, USA ALEX AIKEN, Stanford University, USA

We show that synthesizing recursive functional programs using a class of primitive recursive combinators is both simpler and shows more brechnachtaffs from the literature than previously proposed approaches. Commethed synthesizes paramosphirms, a class of programs that includes the most common recursive programming patterns on algebraic data types. The error of our approach is to applie the synthesis produce minitor buy parts a multi-hole tamplate that fixes the recursive structure, and a search for non-recursive program fragments to fill the template 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

ACM Reference Format:

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 = lappend (Cons h t) l = Cons h (append t l)

This program uses general recursion, that is, the function append is explicitly recursively defined, with ealts to append within its definition. Depending on what other language features are present, unerstricted general recursion is difficult to reason about; for example, proving termination of general 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 patterni where the number of recursive calls is the length of the list argument. A standard (general recursive) definition of fold is:

fold Nil n f = nfold (Cons h t) n f = f h (fold t n f) "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 $g_{\text{Nil}} g_{\text{Cons}} = g_{\text{Nil}}$ para (Cons *h t*) $g_{\text{Nil}} g_{\text{Cons}} = g_{\text{Cons}} h (t, \text{ para } t g_{\text{Nil}} g_{\text{Cons}})$

[...]

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."

Challenges

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?

Practicability

How easy is it for users to express specifications?

Applicability

What range of problems might the system apply to?

Challenges