Component-Based Synthesis of Table Consolidation and Transformation Tasks from Examples *

Yu Feng

University of Texas at Austin, USA
yufeng@cs.utexas.edu

Ruben Martins

University of Texas at Austin, USA
rmartins@cs.utexas.edu

Jacob Van Geffen

University of Texas at Austin, USA
jsv@cs.utexas.edu

Isil Dillig

University of Texas at Austin, USA
isil@cs.utexas.edu

Swarat Chaudhuri

Rice University, USA
swarat@rice.edu

Abstract

This paper presents a novel component-based synthesis algorithm that marries the power of type-directed search with lightweight SMT-based deduction and partial evaluation. Given a set of components together with their *over-approximate* first-order specifications, our method first generates a *program sketch* over a subset of the components and checks its feasibility using an SMT solver. Since a program sketch typically represents *many* concrete programs, the use of SMT-based deduction greatly increases the scalability of the algorithm. Once a feasible program sketch is found, our algorithm completes the sketch in a bottom-up fashion, using partial evaluation to further increase the power of deduction for rejecting partially-filled program sketches. We apply the proposed synthesis methodology for automating a large class of data preparation tasks that commonly arise in data science. We have evaluated our synthesis algorithm on dozens of data wrangling and consolidation tasks obtained from on-line forums, and we show that our approach can automatically solve a large class of problems encountered by R users.

CCS Concepts • **Software and its engineering** → **Programming by example; Automatic programming;** • **Theory of computation** → *Program specifications*

* This work was supported in part by NSF Awards #1453386 and #1162076, and DARPA MUSE Award #8750-14-2-0270.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PLDI'17, June 18–23, 2017, Barcelona, Spain
© 2017 ACM. 978-1-4503-4988-8/17/06...\$15.00
<http://dx.doi.org/10.1145/3062341.3062351>

Keywords Program synthesis, Programming by example, data preparation, Component-based synthesis, SMT-based deduction

1. Introduction

The problem of *program synthesis from examples* has received significant attention from researchers in the last few years. The central objective of this research area is to automate certain classes of programming tasks, either with the goal of helping end-users to “program” or absolving software developers from tedious coding tasks. To accomplish these goals, many program synthesis techniques define the space of relevant programs using a domain-specific language (DSL) and give methods to search the space of DSL programs that are consistent with the user-provided examples. Recent work has shown that such a methodology can be practical in many domains [6, 10, 13, 22, 24, 30].

A particularly interesting version of this problem concerns the synthesis of programs that manipulate *tabular data*. Such programs are especially important in an era where data analytics has gained enormous popularity across a wide range of disciplines, ranging from biology to business to the social sciences. Since raw data is rarely in a form that is immediately amenable to an analytics or visualization task, data scientists typically spend over 80% of their time performing tedious data preparation tasks [7]. Such tasks include consolidating multiple data sources into a single table, reshaping data from one format into another, or adding new rows or columns to an existing table.

While data preparation tasks would seem to be natural targets for synthesis, many such tasks are too complex to be handled by existing techniques. If written in a low-level language, programs implementing these tasks would be simply too large to be discovered by combinatorial search. One way around this difficulty is to describe the relevant computations using a set of predefined library functions, or *compo-*

Figure 1. Overview of our approach

nents, and then synthesize programs that use these high-level primitives. Another advantage of such a *component-based synthesis* approach is its flexibility: Since the reasoning of the synthesizer is not hard-wired to a fixed set of DSL constructs, the underlying algorithm can generate more complex programs as new libraries emerge or as more components are added to its knowledge base.

Unfortunately, a key challenge in developing such a general component-based synthesis algorithm for automating data preparation tasks is scalability: Since many languages (e.g., R) provide a large number of components that are typically used in data preparation, the size of the search space that must be explored by the underlying synthesis algorithm can be very large. Due to this difficulty, prior techniques for automating table transformations (e.g., [17, 37]) focus on narrowly-defined DSLs, such as subsets of the Excel macro language [17] or fragments of SQL [37]. Unfortunately, many common data preparation tasks (e.g., those that involve reshaping tables or require performing nested table joins) fall outside the scope of these previous approaches.

In this paper, we propose a general component-based synthesis algorithm for automating a large class of data preparation tasks. Specifically, our synthesis algorithm is parametrized over a set of components, which can include both higher-order and first-order combinators. The set of components used by the synthesizer can be customized by the user or extended over time as new libraries emerge.

In order to address the scalability challenges that arise from our more general formulation of the problem, we propose a new synthesis algorithm that combines type-directed enumerative search with *lightweight SMT-based deduction* and *partial evaluation*. In our formulation of the synthesis problem, each component C is equipped with a logical, incomplete specification that over-approximates C 's behavior. These specifications are utilized by the synthesizer to perform lightweight SMT-based reasoning, with the goal of rejecting infeasible partial programs. Furthermore, specifications are provided *per component*, so they can be re-used across arbitrarily many synthesis tasks. Since our technique does not depend on hard-coded component-specific reasoning, our approach significantly generalizes prior uses of deduction in example-guided synthesis (e.g., [10]).

Figure 1 shows a schematic illustration of our synthesis algorithm, implemented in a tool called MORPHEUS. To fa-

cilitate effective use of SMT-based deduction, our algorithm decomposes the synthesis task into two separate *sketch generation* and *sketch completion* phases. In particular, a sketch specifies the top-level combinators used in the program, but not their corresponding arguments. Our algorithm uses type-directed enumerative search to lazily explore the space of all possible program sketches and infers a specification of each candidate sketch using the specifications of the underlying components. Hence, once we have a candidate sketch S , we can use an SMT solver to test whether S is consistent with the provided input-output examples. Because a program sketch typically represents *many* concrete programs, the rejection of program sketches using SMT-based reasoning dramatically improves the scalability of the synthesis algorithm.

Once our algorithm finds a feasible program sketch, it then tries to complete it in a bottom-up, type-directed way. In particular, the synthesizer *evaluates* sub-terms of the partial program P to infer a more precise specification for P and again uses SMT-based reasoning with the goal of refuting the partially-completed sketch. Hence, the use of partial evaluation further improves the scalability of the synthesis algorithm by allowing us to refute partial programs obtained during sketch completion.

While the core ideas underlying our algorithm are generally applicable to any component-based synthesizer, we have used these ideas to automate table consolidation and transformation tasks that commonly arise in data science. Specifically, our implementation, MORPHEUS, takes as input a set of source data frames in R, as well as the target data frame that should be generated using the synthesized program. Additionally, the user can also provide a set of components (i.e., library methods), optionally with their corresponding first-order specifications. However, since our implementation already comes with a built-in set of components that are commonly used in data preparation, the user does not need to provide any additional components but can do so if she so desires. Using the ideas outlined above, MORPHEUS then automatically synthesizes an R program that can now be applied to other data frames.

To evaluate our techniques, we have collected a suite of data preparation tasks for the R programming language, drawn from discussions among R users in on-line forums such as Stackoverflow. The “components” in our evaluation are methods provided by two popular R libraries, namely

`tidyr` and `dplyr`, for data tidying and manipulation. Our experiments show that MORPHEUS can successfully synthesize a diverse class of real-world data preparation programs. We also evaluate the performance of MORPHEUS using component specifications of different granularities and demonstrate that SMT-based deduction and partial evaluation are crucial for the scalability of our approach.

To summarize, this paper makes the following key contributions:

- We describe a novel component-based synthesis algorithm that uses SMT-based deduction and partial evaluation to dramatically prune the search space.
- We apply the proposed ideas to automate a diverse class of data wrangling and consolidation tasks that commonly arise in data science.
- We implement these ideas in a tool called MORPHEUS and empirically evaluate our approach in a number of ways. First, we show that MORPHEUS can successfully automate 98% of R-related data preparation tasks collected from on-line forums. Second, we perform a small user study showing that the class of tasks that can be automated by MORPHEUS are difficult even for expert R programmers. Finally, we also show that MORPHEUS can synthesize non-trivial SQL queries and that it performs better than the SQLSYNTHESIZER tool on their own dataset.

2. Motivating Examples

In this section, we illustrate the diversity of data preparation tasks using a few examples collected from Stackoverflow.

Example 1. An R user has the data frame in Figure 2(a), but wants to transform it to the following format [1]:

id	A_2007	B_2007	A_2009	B_2009
1	5	10	5	17
2	3	50	6	17

Even though the user is quite familiar with R libraries for data preparation, she is still not able to perform the desired task. Given this example, MORPHEUS can automatically synthesize the following R program:

```
df1=gather(input, var, val, id, A, B)
df2=unite(df1, yearvar, var, year)
df3=spread(df2, yearvar, val)
```

Observe that this example requires both reshaping the table and appending contents of some cells to column names.

Example 2. Another R user has the data frame from Figure 2(b) and wants to compute, for each source location L , the number and percentage of flights that go to Seattle (SEA) from L [2]. In particular, the output should be as follows:

origin	n	prop
EWR	2	0.6666667
JFK	1	0.3333333

id	year	A	B
1	2007	5	10
2	2009	3	50
1	2007	5	17
2	2009	6	17

(a)

flight	origin	dest
11	EWR	SEA
725	JFK	BQN
495	JFK	SEA
461	LGA	ATL
1696	EWR	ORD
1670	EWR	SEA

(b)

Figure 2. (a) Data frame for Example 1; (b) for Example 2.

MORPHEUS can automatically synthesize the following R program to extract the desired information:

```
df1=filter(input, dest == "SEA")
df2=summarize(group_by(df1, origin), n = n())
df3=mutate(df2, prop = n / sum(n))
```

Observe that this example involves selecting a subset of the data and performing some computation on that subset.

Example 3. A data analyst has the following raw data about the position of vehicles for a driving simulator [3]:

Table 1:

frame	X1	X2	X3
1	0	0	0
2	10	15	0
3	15	10	0

Table 2:

frame	X1	X2	X3
1	0	0	0
2	14.53	12.57	0
3	13.90	14.65	0

Here, Table 1 contains the unique identification number for each vehicle (e.g., 10, 15), with 0 indicating the absence of a vehicle. The column labeled “frame” in Table 1 measures the time step, and the columns “X1”, “X2”, “X3” track which vehicle is closer to the driver. For example, at frame 3, the vehicle with ID 15 is the closest to the driver. Table 2 has a similar structure as Table 1 but contains the speeds of the vehicles instead of their identification number. For example, at frame 3, the speed of the vehicle with ID 15 is 13.90 m/s. The data analyst wants to consolidate these two data frames into a new table with the following shape:

frame	pos	carid	speed
2	X1	10	14.53
3	X2	10	14.65
2	X2	15	12.57
3	X1	15	13.90

Despite looking into R libraries for data preparation, the analyst still cannot figure out how to perform this task and asks for help on Stackoverflow. MORPHEUS can synthesize the following R program to automate this complex task:

```
df1=gather(table1, pos, carid, X1, X2, X3)
df2=gather(table2, pos, speed, X1, X2, X3)
df3=inner_join(df1, df2)
df4=filter(df3, carid != 0)
df5=arrange(df4, carid, frame)
```

3. Problem Formulation

In order to precisely describe our synthesis problem, we first present some definitions that we use throughout the paper.

Definition 1. (Table) A table T is a tuple (r, c, τ, ς) where:

- r, c denote number of rows and columns respectively
- $\tau : \{l_1 : \tau_1, \dots, l_c : \tau_c\}$ denotes the type of T . In particular, each l_i is the name of a column in T and τ_i denotes the type of the value stored in T . We assume that each τ_i is either `num` or `string`.
- ς is a mapping from each cell $(i, j) \in ([0, r) \times [0, c))$ to a value v stored in that cell

Given a table $T = (r, c, \tau, \varsigma)$, we write $T.\text{row}$ and $T.\text{col}$ to denote r and c respectively. We also write $T_{i,j}$ as shorthand for $\varsigma(i, j)$ and $\text{type}(T)$ to represent τ . We refer to all record types $\{l_1 : \tau_1, \dots, l_c : \tau_c\}$ as type `tbl`. In addition, tables with only one row are referred to as being of type `row`.

Definition 2. (Component) A component $\mathcal{X}$ is a triple (f, τ, ϕ) where f is a string denoting $\mathcal{X}$'s name, τ is the type signature (see Figure 3), and ϕ is a first-order formula that specifies $\mathcal{X}$'s input-output behavior.

Given a component $\mathcal{X} = (f, \tau, \phi)$, the specification ϕ is over the vocabulary $x_1, \dots, x_n, y$, where x_i denotes $\mathcal{X}$'s i 'th argument and y denotes $\mathcal{X}$'s return value. Note that specification ϕ does not need to *precisely* capture $\mathcal{X}$'s input-output behavior; it only needs to be an *over-approximation*. Thus, *true* is always a valid specification for any component.

With slight abuse of notation, we sometimes write $\mathcal{X}(\dots)$ to mean $f(\dots)$ whenever $\mathcal{X} = (f, \tau, \phi)$. Also, given a component $\mathcal{X}$ and arguments $c_1, \dots, c_n$, we write $\llbracket \mathcal{X}(c_1, \dots, c_n) \rrbracket$ to denote the result of evaluating $\mathcal{X}$ on arguments $c_1, \dots, c_n$.

Definition 3. (Problem specification) The specification for a synthesis problem is a pair $(\mathcal{E}, \Lambda)$ where:

- $\mathcal{E}$ is an input-output example $(\vec{T}_{\text{in}}, T_{\text{out}})$ such that $\vec{T}_{\text{in}}$ denotes a list of input tables, and T_{out} is the output table,
- $\Lambda = (\Lambda_T \cup \Lambda_v)$ is a set of components, where Λ_T, Λ_v denote table transformers and value transformers respectively. We assume that Λ_T includes higher-order functions, but Λ_v consists of first-order operators.

Given an input-output example $\mathcal{E} = (\vec{T}_{\text{in}}, T_{\text{out}})$, we write $\mathcal{E}_{\text{in}}, \mathcal{E}_{\text{out}}$ to denote $\vec{T}_{\text{in}}, T_{\text{out}}$ respectively. Also, we classify components Λ into two disjoint classes Λ_T and Λ_v , where Λ_T denotes *table transformer* components that take at least one table as an argument and return a table. Components of all other types are *value transformers* Λ_v . While table transformers can be higher-order combinators, value transformers are always first-order. In the rest of the paper, we assume that table transformers only take tables and first-order functions (constructed using constants and components in Λ_v) as arguments.

Cell type γ	$:=$ num string
Primitive type β	$:=$ γ bool cols
Table type <code>tbl</code>	$:=$ $\{l_1 : \gamma_1, \dots, l_n : \gamma_n\}$ (row $<:$ tbl)
Type τ	$:=$ β tbl $\tau_1 \rightarrow \tau_2$ $\tau_1 \times \tau_2$

Figure 3. Types used in components; `cols` represents a list of strings where each string is a column name in some table.

Example 4. Consider the selection operator σ from relational algebra, which takes a table and a predicate and returns a table. In our terminology, such a component is a table transformer. In contrast, an aggregate function such as *sum* that takes a list of values and returns their sum is a value transformer. Similarly, the boolean operator $\geq$ is also a value transformer.

Definition 4. (Synthesis problem) Given specification $(\mathcal{E}, \Lambda)$ where $\mathcal{E} = (\vec{T}_{\text{in}}, T_{\text{out}})$, the synthesis problem is to infer a program $\lambda \vec{x}.e$ such that (a) e is a well-typed expression over components in Λ , and (b) $(\lambda \vec{x}.e) \vec{T}_{\text{in}} = T_{\text{out}}$.

4. Hypotheses as Refinement Trees

Before we can describe our synthesis algorithm, we first introduce *hypotheses* that represent partial programs with unknown expressions (i.e., holes). More formally, hypotheses $\mathcal{H}$ are defined by the grammar presented in Figure 5. In the simplest form, a hypothesis $(?_i : \tau)$ represents an unknown expression of type τ . More complicated hypotheses are constructed using table transformation components $\mathcal{X} \in \Lambda_T$. In particular, if $\mathcal{X} = (f, \tau, \phi) \in \Lambda_T$, a hypothesis of the form $?_i^{\mathcal{X}}(\mathcal{H}_1, \dots, \mathcal{H}_n)$ represents an expression $f(e_1, \dots, e_n)$.

During the course of our synthesis algorithm, we will progressively fill the holes in the hypothesis with concrete expressions. For this reason, we also allow hypotheses of the form $(?_i : \tau) @ \mathcal{Q}$ where *qualifier* $\mathcal{Q}$ specifies the term that is used to fill hole $?_i$. Specifically, if $?_i$ is of type `tbl`, then its corresponding qualifier has the form (x, T) , which means that $?_i$ is instantiated with input variable x , which is in turn bound to table T in the input-output example provided by the user. On the other hand, if $?_i$ is of type $(\tau_1 \times \dots \times \tau_n) \rightarrow \tau$, then the qualifier must be a first-order function $\lambda y_1. \dots y_n. t$ constructed using components Λ_v .¹

Our synthesis algorithm starts with the most general hypothesis and progressively makes it more specific. Therefore, we now define what it means to *refine* a hypothesis:

Definition 5. (Hypothesis refinement) Given two hypotheses $\mathcal{H}, \mathcal{H}'$, we say that $\mathcal{H}'$ is a refinement of $\mathcal{H}$ if it can be obtained by replacing some subterm $?_i : \tau$ of $\mathcal{H}$ by $?_i^{\mathcal{X}}(\mathcal{H}_1, \dots, \mathcal{H}_n)$ where $\mathcal{X} = (f, \tau' \rightarrow \tau, \phi) \in \Lambda_T$.

In other words, a hypothesis $\mathcal{H}'$ refines another hypothesis $\mathcal{H}$ if it makes it more constrained.

¹ We view constants as a special case of first-order functions.

$$\begin{aligned}
\llbracket (?_i : \tau) \rrbracket_{\partial} &= ?_i & \llbracket (?_i : \tau) @ (x, T) \rrbracket_{\partial} &= T & \llbracket (?_i : \tau) @ t \rrbracket_{\partial} &= t \\
\llbracket ?_i^{\chi}(\mathcal{H}_1, \dots, \mathcal{H}_n) \rrbracket_{\partial} &= \begin{cases} \mathcal{X}(\llbracket \mathcal{H}_1 \rrbracket_{\partial}, \dots, \llbracket \mathcal{H}_n \rrbracket_{\partial}) & \text{if } \exists i \in [1, n]. \text{ PARTIAL}(\llbracket \mathcal{H}_i \rrbracket_{\partial}) \\ \llbracket \mathcal{X}(\llbracket \mathcal{H}_1 \rrbracket_{\partial}, \dots, \llbracket \mathcal{H}_n \rrbracket_{\partial}) \rrbracket_{\partial} & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 4. Partial evaluation of hypothesis. We write $\text{PARTIAL}(\llbracket \mathcal{H} \rrbracket_{\partial})$ if $\llbracket \mathcal{H} \rrbracket_{\partial}$ contains at least one question mark.

Term t $:= \text{const} \mid y_i \mid \mathcal{X}(t_1, \dots, t_n) \ (\mathcal{X} \in \Lambda_v)$
 Qualifier $\mathcal{Q}$ $:= (x, T) \mid \lambda y_1, \dots, y_n. t$
 Hypothesis $\mathcal{H}$ $:= (?_i : \tau) \mid (?_i : \tau) @ \mathcal{Q}$
 $\quad \mid ?_i^{\chi}(\mathcal{H}_1, \dots, \mathcal{H}_n) \ (\mathcal{X} \in \Lambda_T)$

Figure 5. Context-free grammar for hypotheses

Figure 6. Representing hypotheses as refinement trees

Figure 7. A sketch (left) and a complete program (right)

Example 5. The hypothesis $\mathcal{H}_1 = ?_0^{\sigma} (?_1 : \text{tbl}, ?_2 : \text{row} \rightarrow \text{bool})$ is a refinement of $\mathcal{H}_0 = ?_0 : \text{tbl}$ because $\mathcal{H}_1$ is more specific than $\mathcal{H}_0$. In particular, $\mathcal{H}_0$ represents any arbitrary expression of type tbl , whereas $\mathcal{H}_1$ represents expressions whose top-level construct is a selection.

Since our synthesis algorithm starts with the hypothesis $?_0 : \text{tbl}$ and iteratively refines it, we will represent hypotheses using *refinement trees* [24]. Effectively, a refinement tree corresponds to the *abstract syntax tree (AST)* for the hypotheses from Figure 5. In particular, note that internal nodes labeled $?_i^{\chi}$ of a refinement tree represent hypotheses whose top-level construct is χ . If an internal node $?_i^{\chi}$ has children labeled with unknowns $?_j, \dots, ?_{j+n}$, this means that hypothesis $?_i$ was refined to $\chi(?_j, \dots, ?_{j+n})$. Intuitively, a refinement tree captures the *history* of refinements that occur as we search for the desired program.

Example 6. Consider the refinement tree from Figure 6, and suppose that π, σ denote the standard projection and selection operators in relational algebra. This refinement tree represents the partial program $\pi(\sigma(?_1, ?_2))$. The refinement tree also captures the search history in our synthesis algorithm. Specifically, it shows that our initial hypothesis was $?_0$, which then got refined to $\pi(?_1, ?_2)$, which in turn was refined to $\pi(\sigma(?_3, ?_4), ?_2)$.

T ₁				T ₂			
id	name	age	GPA	id	name	age	GPA
1	Alice	8	4.0	2	Bob	18	3.2
2	Bob	18	3.2	3	Tom	12	3.0
3	Tom	12	3.0				

Figure 8. Tables for Example 8

As mentioned in Section 1, our approach decomposes the synthesis task into two separate *sketch generation* and *sketch completion* phases. We define a *sketch* to be a special kind of hypothesis where there are no unknowns of type tbl .

Definition 6. (Sketch) A sketch is a special form of hypothesis where all leaf nodes of type tbl have a corresponding qualifier of the form (x, T) .

In other words, a sketch completely specifies the table transformers used in the target program, but the first-order functions supplied as arguments to the table transformers are yet to be determined.

Example 7. Consider the refinement tree from Figure 6. This hypothesis is not a sketch because there is a leaf node (namely $?_3$) of type tbl that does not have a corresponding qualifier. On the other hand, the refinement tree shown in Figure 7 (left) is a sketch and corresponds to the partial program $\pi(x_1, ?)$ where $?$ is a list of column names. Furthermore, this sketch states that variable x_1 corresponds to table T from the input-output example.

Definition 7. (Complete program) A complete program is a hypothesis where all leaf nodes are of the form $(?_i : \tau) @ \mathcal{Q}$.

In other words, a complete program fully specifies the expression represented by each $?$ in the hypothesis. For instance, a hypothesis that represents a complete program is shown in Figure 7 (right) and represents the relational algebra term $\lambda x_1. \pi_{\text{name, year}}(x_1)$.

As mentioned in Section 1, our synthesis procedure relies on performing partial evaluation. Hence, we define a function $\llbracket \mathcal{H} \rrbracket_{\partial}$, shown in Figure 4, for partially evaluating hypothesis $\mathcal{H}$. Observe that, if $\mathcal{H}$ is a complete program, then $\llbracket \mathcal{H} \rrbracket_{\partial}$ evaluates to a concrete table. Otherwise, $\llbracket \mathcal{H} \rrbracket_{\partial}$ returns a partially evaluated hypothesis. We write $\text{PARTIAL}(\llbracket \mathcal{H} \rrbracket_{\partial})$ if $\llbracket \mathcal{H} \rrbracket_{\partial}$ does not evaluate to a concrete term (i.e., contains question marks).

Example 8. Consider hypothesis $\mathcal{H}$ on the left-hand side of Figure 9, where T_1 is Table 1 from Figure 8. The refinement tree on the right-hand-side of Figure 9 shows the result of partially evaluating $\mathcal{H}$, where T_2 is Table 2 from Figure 8.

Figure 9. Partial evaluation on hypothesis from Figure 6; $\text{age} > 8$ stands for $?_4 : \text{row} \rightarrow \text{bool} @ \lambda x. (x.\text{age} > 8)$.

Figure 10. Illustration of the top-level synthesis algorithm

5. Synthesis Algorithm

In this section, we describe the high-level structure of our synthesis algorithm, leaving the discussion of SMT-based deduction and sketch completion to the next two sections.

As illustrated schematically in Figure 10, our synthesis algorithm maintains a priority queue of hypotheses, which are either converted into a sketch or refined to a more specific hypothesis during each iteration. Specifically, the synthesis procedure picks the most promising hypothesis $\mathcal{H}$ according to some heuristic cost metric (explained in Section 8) and asks the deduction engine if $\mathcal{H}$ can be successfully converted into a sketch. If the deduction engine refutes this conjecture, we then discard $\mathcal{H}$ but add all possible (one-level) refinements of $\mathcal{H}$ into the worklist. Otherwise, we convert hypothesis $\mathcal{H}$ into a sketch $\mathcal{S}$ and try to complete it using the *sketch completion engine*.

Algorithm 1 describes our top-level synthesis algorithm in more detail. Given an example $\mathcal{E}$ and a set of components Λ , SYNTHESIZE either returns a complete program that satisfies $\mathcal{E}$ or yields $\perp$, meaning that no such program exists.

Internally, the SYNTHESIZE procedure maintains a priority queue W of all hypotheses. Initially, the only hypothesis

$$\frac{T_j \in T_{in}}{\mathcal{H} = (?_i : \text{tbl})} \quad (1)$$

$$\frac{}{\mathcal{H} @ (x_j, T_j) \in \text{Sketches}(\mathcal{H}, \vec{T}_{in})}$$

$$\frac{\mathcal{H} = ?_i : \tau_i}{\tau_i \neq \text{tbl}} \quad (2)$$

$$\frac{}{\mathcal{H} \in \text{Sketches}(\mathcal{H}, \vec{T}_{in})}$$

$$\frac{\mathcal{H} = ?_i^{\mathcal{X}}(\mathcal{H}_1, \dots, \mathcal{H}_n)}{\mathcal{H}'_i \in \text{Sketches}(\mathcal{H}_i, \vec{T}_{in})} \quad (3)$$

$$\frac{}{?_i^{\mathcal{X}}(\mathcal{H}'_1, \dots, \mathcal{H}'_n) \in \text{Sketches}(\mathcal{H}, \vec{T}_{in})}$$

Figure 11. Converting a hypothesis into a sketch.

Algorithm 1 Synthesis Algorithm

```

1: procedure SYNTHESIZE( $\mathcal{E}, \Lambda$ )
2:   input: Input-output example  $\mathcal{E}$  and components  $\Lambda$ 
3:   output: Synthesized program or  $\perp$  if failure
4:    $W := \{?_0 : \text{tbl}\}$  ▷ Init worklist
5:   while  $W \neq \emptyset$  do
6:     choose  $\mathcal{H} \in W$ ;
7:      $W := W \setminus \{\mathcal{H}\}$ 
8:     if DEDUCE( $\mathcal{H}, \mathcal{E}$ ) =  $\perp$  then ▷ Contradiction
9:       goto refine;
10:    ▷ No contradiction
11:    for  $\mathcal{S} \in \text{SKETCHES}(\mathcal{H}, \mathcal{E}_{in})$  do
12:       $\mathcal{P} := \text{FILLSKETCH}(\mathcal{S}, \mathcal{E})$ 
13:      for  $p \in \mathcal{P}$  do
14:        if CHECK( $p, \mathcal{E}$ ) then return  $p$ 
15:    refine: ▷ Hypothesis refinement
16:    for  $\mathcal{X} \in \Lambda_T, (?_i : \text{tbl}) \in \text{LEAVES}(\mathcal{H})$  do
17:       $\mathcal{H}' := \mathcal{H}[?_i^{\mathcal{X}}(?_j : \tau)/?_i]$ 
18:       $W := W \cup \mathcal{H}'$ 
19:  return  $\perp$ 

```

in W is $?_0$, which represents any possible program. In each iteration of the while loop (lines 5–18), we pick a hypothesis $\mathcal{H}$ from W and invoke the DEDUCE procedure (explained later) to check if $\mathcal{H}$ can be directly converted into a sketch by filling holes of type `tbl` with the input variables. Note that our deduction procedure is sound but, in general, not complete: Since component specifications are over-approximate, the deduction procedure can return $\top$ (i.e., true) even though no valid completion of the sketch exists. However, DEDUCE returns $\perp$ only when the current hypothesis requires further refinement. Hence, the use of deduction does not lead to a loss of completeness in our overall synthesis approach.

If DEDUCE does not find a conflict, we then convert the current hypothesis $\mathcal{H}$ into a set of possible sketches (line 11). The function SKETCHES used at line 11 is presented using inference rules in Figure 11. Effectively, we convert hypothesis $\mathcal{H}$ into a sketch by replacing each hole of type `tbl` with one of the input variables x_j , which corresponds to table T_j in the input-output example.

Now, given a candidate sketch $\mathcal{S}$, we try to complete it using the call to FILLSKETCH at line 12 (explained in Section 7). FILLSKETCH returns a *set* of complete programs $\mathcal{P}$ such that each $p \in \mathcal{P}$ is valid with respect to our deduction procedure. However, as our deduction procedure is incomplete, p may not satisfy the input-output examples. Hence, we only return p as a solution if p satisfies $\mathcal{E}$ (line 14).

Lines 16–18 of Algorithm 1 perform *hypothesis refinement*. The idea behind hypothesis refinement is to replace one of the holes of type `tbl` in $\mathcal{H}$ with a component from Λ_T , thereby obtaining a more specific hypothesis. Each of

$$\begin{aligned}
\Phi(\mathcal{H}_i) &= \alpha(\llbracket \mathcal{H}_i \rrbracket_{\partial})[?_i/x] \text{ if } \neg \text{PARTIAL}(\llbracket \mathcal{H}_i \rrbracket_{\partial}) \\
\Phi(\mathcal{H}_i) &= \top \text{ else if ISLEAF}(\mathcal{H}_i) \\
\Phi(?_0^{\mathcal{X}}(\mathcal{H}_1, \dots, \mathcal{H}_n)) &= \bigwedge_{1 \leq i \leq n} \Phi(\mathcal{H}_i) \wedge \phi_{\mathcal{X}}[?_0/y, \vec{?}_i/\vec{x}_i]
\end{aligned}$$

Figure 12. Constraint generation for hypotheses. $?_i$ denotes the root variable of $\mathcal{H}_i$ and the specification of $\mathcal{X}$ is $\phi_{\mathcal{X}}$. Function α generates an SMT formula describing its input table.

the refined hypotheses is added to the worklist and possibly converted into a sketch in future iterations.

6. SMT-based Deduction

In the previous section, we described the structure of the synthesis algorithm, but did not yet explain the underlying deductive reasoning engine. The key idea here is to generate an SMT formula that corresponds to the specification of the current sketch and to check whether the input-output example satisfies this specification.

Component specifications. We use the specifications of individual components to derive the overall specification for a given hypothesis. As mentioned earlier, these specifications need not be precise and can, in general, over-approximate the behavior of the components. For instance, Table 1 shows sample specifications for a subset of methods from two popular R libraries. Note that these sample specifications do not fully capture the behavior of each component and only describe the relationship between the number of rows and columns in the input and output tables.² For example, consider the `filter` function from the `dplyr` library for selecting a subset of the rows that satisfy a given predicate in the data frame. The specification of `filter`, which is effectively the selection operator σ from relational algebra, is given by:

$$T_{out}.row < T_{in}.row \wedge T_{out}.col = T_{in}.col$$

In other words, this specification expresses that the table obtained after applying the `filter` function contains fewer rows but the same number of columns as the input table.³

Generating specification for hypothesis. Given a hypothesis $\mathcal{H}$, we need to generate the specification for $\mathcal{H}$ using the specifications of the individual components used in $\mathcal{H}$. Towards this goal, the function $\Phi(\mathcal{H})$ defined in Figure 12 returns the specification of hypothesis $\mathcal{H}$.

In the simplest case, $\mathcal{H}_i$ corresponds to a complete program (line 1 of Figure 12)⁴. In this case, we evaluate the hy-

²The actual specifications used in our implementation are slightly more involved. In Section 9, we compare the performance of MORPHEUS using two different specifications.

³In principle, the number of rows may be unchanged if the predicate does not match any row. However, we need not consider this case since there is a simpler program without `filter` that satisfies the example.

⁴Recall that the DEDUCE procedure will also be used during sketch completion. While $\mathcal{H}$ can never be a complete program when called from line 8

pothesis to a table T and obtain $\Phi(\mathcal{H}_i)$ as the “abstraction” of T . In particular, the *abstraction function* α used in Figure 12 takes as input a concrete table T and returns a constraint describing that table. In general, the definition of the abstraction function α depends on the granularity of the component specifications. For instance, if our component specifications only refer to the number of rows and columns, then a suitable abstraction function for an $m \times n$ table would yield $x.row = m \wedge x.col = n$. In general, we assume variable x is used to describe the input table of α .

Let us now consider the second case in Figure 12 where $\mathcal{H}_i$ is a leaf, but not a complete program. In this case, since we do not have any information about what $\mathcal{H}_i$ represents, we return $\top$ (i.e., *true*) as the specification.

Finally, let us consider the case where the hypothesis is of the form $?_0^{\mathcal{X}}(\mathcal{H}_1, \dots, \mathcal{H}_n)$. In this case, we first recursively infer the specifications of sub-hypotheses $\mathcal{H}_1, \dots, \mathcal{H}_n$. Now suppose that the specification of $\mathcal{X}$ is given by $\phi_{\mathcal{X}}(\vec{x}, y)$, where $\vec{x}$ and y denote $\mathcal{X}$ ’s inputs and output respectively. If the root variable of each hypothesis $\mathcal{H}_i$ is given by $?_i$, then the specification for the overall hypothesis is obtained as:

$$\bigwedge_{1 \leq i \leq n} \Phi(\mathcal{H}_i) \wedge \phi_{\mathcal{X}}[?_0/y, \vec{?}_i/\vec{x}_i]$$

Example 9. Consider hypothesis $\mathcal{H}$ from Figure 6, and suppose that the specifications for relational algebra operators π and σ are the same as `select` and `filter` from Table 1 respectively. Then, $\Phi(\mathcal{H})$ corresponds to the following Presburger arithmetic formula:

$$\begin{aligned}
&?_1.row < ?_3.row \wedge ?_1.col = ?_3.col \wedge \\
&?_0.row = ?_1.row \wedge ?_0.col < ?_1.col
\end{aligned}$$

Here, $?_3, ?_0$ denote the input and output tables respectively, and $?_1$ is the intermediate table obtained after selection.

Deduction using SMT. Algorithm 2 presents our deduction algorithm using the constraint generation function Φ defined in Figure 12. Given a hypothesis $\mathcal{H}$ and input-output example $\mathcal{E}$, DEDUCE returns $\perp$ if $\mathcal{H}$ does not correspond to a valid sketch. In other words, $\text{DEDUCE}(\mathcal{H}, \mathcal{E}) = \perp$ means that we cannot obtain a program that satisfies the input-output examples by replacing holes with inputs.

As shown in Algorithm 2, the DEDUCE procedure generates a constraint ψ and checks its satisfiability using an SMT solver. If ψ is unsatisfiable, hypothesis $\mathcal{H}$ cannot be unified with the input-output example and can therefore be rejected.

Let us now consider the construction of SMT formula ψ in Algorithm 2. First, given a hypothesis $\mathcal{H}$, the corresponding sketch must map each of the unknowns of type `tbl` to one of the arguments. Hence, the constraint φ_{in} generated at line 5 indicates that each leaf with label $?_j$ corresponds to some argument x_i . Similarly, φ_{out} expresses that

of the SYNTHESIZE procedure (Algorithm 1), it can be a complete program when DEDUCE is invoked through the sketch completion engine.

Lib	Component	Description	Specification
tidyr	spread	Spread a key-value pair across multiple columns.	$T_{out}.row \leq T_{in}.row$ $T_{out}.col \geq T_{in}.col$
	gather	Takes multiple columns and collapses into key-value pairs, duplicating all other columns as needed.	$T_{out}.row \geq T_{in}.row$ $T_{out}.col \leq T_{in}.col$
dplyr	select	Project a subset of columns in a data frame.	$T_{out}.row = T_{in}.row$ $T_{out}.col < T_{in}.col$
	filter	Select a subset of rows in a data frame.	$T_{out}.row < T_{in}.row$ $T_{out}.col = T_{in}.col$

Table 1. Sample specifications of a few components

Algorithm 2 SMT-based Deduction Algorithm

```

1: procedure DEDUCE( $\mathcal{H}, \mathcal{E}$ )
2:   input: Hypothesis  $\mathcal{H}$ , input-output example  $\mathcal{E}$ 
3:   output:  $\perp$  if cannot be unified with  $\mathcal{E}$ ;  $\top$  otherwise
4:    $\mathcal{S} := \{?_j \mid ?_j : \text{tbl} \in \text{LEAVES}(\mathcal{H})\}$ 
5:    $\varphi_{in} := \bigwedge_{?_j \in \mathcal{S}} \bigvee_{1 \leq i \leq |\mathcal{E}_{in}|} (?_j = x_i)$ 
6:    $\varphi_{out} := (y = \text{ROOTVAR}(\mathcal{H}))$ 
7:    $\psi := \left( \bigwedge_{T_i \in \mathcal{E}_{in}} (\alpha(T_i)[x_i/x]) \wedge \alpha(T_{out})[y/x] \right)$ 
8:   return SAT( $\psi$ )

```

the root variable of hypothesis $\mathcal{H}$ must correspond to the return value y of the synthesized program. Hence, the constraint $\Phi(\mathcal{H}) \wedge \varphi_{in} \wedge \varphi_{out}$ expresses the specification of the sketch in terms of variables $x_1, \dots, x_n, y$.

Now, to check if $\mathcal{H}$ is unifiable with example $\mathcal{E}$, we must also generate constraints that describe each table T_{in}^i in terms of x_i and T_{out} in terms of y . Recall from earlier that the abstraction function $\alpha(T)$ generates an SMT formula describing T in terms of variable x . Hence, the constraint

$$\bigwedge_{T_i \in \mathcal{E}_{in}} (\alpha(T_i)[x_i/x]) \wedge \alpha(T_{out})[y/x]$$

expresses that each T_{in}^i must correspond to x_i and T_{out} must correspond to variable y . Thus, the unsatisfiability of formula ψ at line 7 indicates that hypothesis $\mathcal{H}$ can be rejected.

Example 10. Consider the hypothesis from Figure 6, and suppose that the input and output tables are T_1 and T_2 from Figure 8 respectively. The DEDUCE procedure from Algorithm 2 generates the following constraint ψ :

$$\begin{aligned}
& ?_1.\text{row} < ?_3.\text{row} \wedge ?_1.\text{col} = ?_3.\text{col} \wedge ?_0.\text{row} = ?_1.\text{row} \\
& \wedge ?_0.\text{col} < ?_1.\text{col} \wedge x_1 = ?_3 \wedge y = ?_0 \wedge \\
& x_1.\text{row} = 3 \wedge x_1.\text{col} = 4 \wedge y.\text{row} = 2 \wedge y.\text{col} = 4
\end{aligned}$$

Observe that $\Phi(\mathcal{H}) \wedge \varphi_{in} \wedge \varphi_{out}$ implies $y.\text{col} < x_1.\text{col}$, indicating that the output table should have fewer columns than the input table. Since we have $x_1.\text{col} = y.\text{col}$, constraint ψ is unsatisfiable, allowing us to reject the hypothesis.

$$\begin{aligned}
& \frac{\text{type}(T) = \{l_1 : \tau_1, \dots, l_n : \tau_n\} \quad c = [l_i \mid i \in C_i] \text{ for } C_i \in \mathcal{P}([1, n])}{\Gamma \vdash c \in \Omega(\text{cols}, T)} \quad (\text{Cols}) \\
& \frac{c \in T, \text{type}(c) = \tau \quad \tau \in \{\text{num}, \text{string}\}}{\Gamma \vdash c \in \Omega(\tau, T)} \quad (\text{Const}) \\
& \frac{\Gamma \vdash x : \tau}{\Gamma \vdash x \in \Omega(\tau, T)} \quad (\text{Var}) \\
& \frac{\Gamma \vdash t_1 \in \Omega(\tau_1, T) \quad \Gamma \vdash t_2 \in \Omega(\tau_2, T)}{\Gamma \vdash (t_1, t_2) \in \Omega(\tau_1 \times \tau_2, T)} \quad (\text{Tuple}) \\
& \frac{(f, \tau' \rightarrow \tau, \phi) \in \Lambda_v \quad \Gamma \vdash t \in \Omega(\tau', T)}{\Gamma \vdash f(t) \in \Omega(\tau, T)} \quad (\text{App}) \\
& \frac{\tau = (\tau_1 \times \dots \times \tau_n \rightarrow \tau') \quad \Gamma' = \Gamma \cup \{x_1 : \tau_1, \dots, x_n : \tau_n\} \quad \Gamma' \vdash t \in \Omega(\tau', T)}{\Gamma \vdash (\lambda x_1, \dots, x_n. t) \in \Omega(\tau, T)} \quad (\text{Lambda})
\end{aligned}$$

Figure 13. Table-driven type inhabitation rules.

7. Sketch Completion

The goal of sketch completion is to fill the remaining holes in the hypothesis with first-order functions constructed using components in Λ_v . For instance, consider the sketch $\pi(\sigma(x, ?_1), ?_2)$ where π, σ are the familiar projection and selection operators from relational algebra. Now, in order to fill hole $?_1$, we need to know the columns in table x . Similarly, in order to fill hole $?_2$, we need to know the columns in the intermediate table obtained using selection.

As this example illustrates, the vocabulary of first-order functions that can be supplied as arguments to table transformers often depends on the shapes (i.e., schemas) of the other arguments of type `tbl`. For this reason, our sketch completion algorithm synthesizes the program *bottom-up*, evaluating terms of type `tbl` before synthesizing the other arguments. Furthermore, as discussed in Section 1, the completion of program sketches in a bottom-up manner allows us to perform partial evaluation, which in turn increases the effectiveness of the deductive reasoning engine.

$$\frac{\begin{array}{l} \mathcal{S} = (?_i : \tau_i) \\ t \in \Omega(\tau_i, \mathbf{T}, \emptyset) \\ \text{DEDUCE}(\mathcal{S}_f[\mathcal{S}@t/\mathcal{S}], \mathcal{E}) \neq \perp \end{array}}{\mathcal{S}@t \in \mathcal{C}_v(\mathcal{S}, \mathcal{S}_f, \mathcal{E}, \mathbf{T})} \quad (1)$$

$$\frac{\mathcal{S} = (?_i, \text{tbl})@(x, \mathbf{T})}{(\mathcal{S}, \mathbf{T}) \in \mathcal{C}_\tau(\mathcal{S}, \mathcal{S}_f, \mathcal{E})} \quad (2)$$

$$\frac{\begin{array}{l} \mathcal{S} = ?_i^x(\vec{\mathcal{H}} : \text{tbl}, \vec{\mathcal{H}}' : \tau) \quad (\tau \neq \text{tbl}) \\ (\mathcal{P}_j, \mathbf{T}_j) \in \mathcal{C}_\tau(\mathcal{H}_j, \mathcal{S}_f, \mathcal{E}) \\ \mathcal{P}'_j \in \mathcal{C}_v(\mathcal{H}'_j, \mathcal{S}_f[\vec{\mathcal{P}}/\vec{\mathcal{H}}], \mathcal{E}, \mathbf{T}_1 \times \dots \times \mathbf{T}_n) \\ \text{DEDUCE}(\mathcal{S}_f[\vec{\mathcal{P}}/\vec{\mathcal{H}}, \vec{\mathcal{P}}'/\vec{\mathcal{H}}'], \mathcal{E}) \neq \perp \\ \mathcal{P}^* = \mathcal{S}[\vec{\mathcal{P}}/\vec{\mathcal{H}}, \vec{\mathcal{P}}'/\vec{\mathcal{H}}'] \end{array}}{(\mathcal{P}^*, \llbracket \mathcal{P}^* \rrbracket_\partial) \in \mathcal{C}_\tau(\mathcal{S}, \mathcal{S}_f, \mathcal{E})} \quad (3)$$

$$\frac{(\mathcal{P}, \mathbf{T}) \in \mathcal{C}_\tau(\mathcal{S}, \mathcal{S}, \mathcal{E})}{\mathcal{P} \in \text{FILLSKETCH}(\mathcal{S}, \mathcal{E})} \quad (4)$$

Figure 14. Sketch completion rules.

Table-driven type inhabitation. At a high level, our sketch completion procedure is type-directed and synthesizes an argument of type τ by enumerating all inhabitants of τ . However, as argued earlier, the valid inhabitants of type τ are determined by a particular table. Hence, we consider the table-driven variant of the standard type inhabitation problem: That is, given a type τ and a concrete table $\mathbf{T}$, what are all valid inhabitants of τ with respect to the universe of constants used in $\mathbf{T}$?

We formalize this variant of the type inhabitation problem using the inference rules shown in Figure 13. Specifically, these rules derive judgments of the form $\Gamma \vdash t \in \Omega(\tau, \mathbf{T})$ where Γ is a type environment mapping variables to types. The meaning of this judgment is that, under type environment Γ , term t is a valid inhabitant of type τ with respect to table $\mathbf{T}$. Observe that we need the type environment Γ due to the presence of function types: That is, given a function type $\tau_1 \rightarrow \tau_2$, we need Γ to enumerate valid inhabitants of τ_2 . Since the typing rules from Figure 13 resemble those for the simply-typed lambda calculus, we do not explain them in detail. The main difference is that constants of type `cols` are drawn from lists of column names from the table schema, and constants of type `num` and `string` are drawn from values in the table.

Example 11. Consider table $\mathbf{T}_1$ from Figure 8 and the type environment $\Gamma : \{x \mapsto \text{string}\}$. Assuming $\text{eq} : \text{string} \times \text{string} \rightarrow \text{bool}$ is a component in Λ_v , we have $\text{eq}(x, \text{"Alice"}) \in \Omega(\text{bool}, \mathbf{T}_1)$ using the `App`, `Const`, `Var` rules. Similarly, $\lambda x. \text{eq}(x, \text{"Bob"})$ is also a valid inhabitant of $\text{string} \rightarrow \text{bool}$ with respect to $\mathbf{T}_1$.

Sketch completion algorithm. Our sketch completion procedure is described using the inference rules shown in Figure 14. As mentioned previously, the algorithm is bottom-up and first synthesizes all arguments of type `tbl` before syn-

$\mathbf{T}_3$			$\mathbf{T}_4$			
id	name	age	id	name	age	GPA
2	Bob	18	2	Bob	18	3.2
3	Tom	12				

Figure 15. Tables for Example 12

thesizing other arguments. Given sketch $\mathcal{S}$ and example $\mathcal{E}$, $\text{FILLSKETCH}(\mathcal{S}, \mathcal{E})$ returns a set of hypotheses representing complete well-typed programs that are valid with respect to our deduction system.

The first rule in Figure 14 corresponds to a base case of the FILLSKETCH procedure and is used for completing hypotheses that are *not* of type `tbl`. Here, $\mathcal{S}$ represents a sub-part of the sketch that we want to complete, $\mathbf{T}$ is the table that should be used in completing $\mathcal{S}$, and $\mathcal{S}_f$ is the full sketch. Since $\mathcal{S}$ represents an unknown expression of type τ_i , we use the type inhabitation rules from Figure 13 to find a well-typed instantiation t of τ_i with respect to table $\mathbf{T}$. Given completion t of $?_i$, the full sketch now becomes $\mathcal{S}_f[\mathcal{S}@t/\mathcal{S}]$, and we use the deduction system to check whether the new hypothesis is valid. Since our deduction procedure uses partial evaluation, we may now be able to obtain a concrete table for some part of the sketch, thereby enhancing the power of deductive reasoning.

The second rule from Figure 14 is also a base case of the FILLSKETCH procedure. Since any leaf $?_i$ of type `tbl` is already bound to some input variable x in the sketch, there is nothing to complete; hence, we just return $\mathcal{S}$ itself.

Rule (3) corresponds to the recursive step of the FILLSKETCH procedure and is used to complete a sketch with top-most component χ . Specifically, consider a sketch of the form $?_i^x(\vec{\mathcal{H}}, \vec{\mathcal{H}}')$ where $\vec{\mathcal{H}}$ denotes arguments of type `tbl` and $\vec{\mathcal{H}}'$ represents first-order functions. Since the vocabulary of $\vec{\mathcal{H}}'$ depends on the completion of $\vec{\mathcal{H}}$ (as explained earlier), we first recursively synthesize $\vec{\mathcal{H}}$ and obtain a set of complete programs $\vec{\mathcal{P}}$, together with their partial evaluation $\mathbf{T}_1, \dots, \mathbf{T}_n$. Now, observe that each $\mathcal{H}'_j \in \vec{\mathcal{H}}'$ can refer to any of the columns in $\mathbf{T}_1 \times \dots \times \mathbf{T}_n$; hence we recursively synthesize the remaining arguments $\vec{\mathcal{H}}'$ using table $\mathbf{T}_1 \times \dots \times \mathbf{T}_n$. Now, suppose that the hypotheses $\vec{\mathcal{H}}$ and $\vec{\mathcal{H}}'$ are completed using terms $\vec{\mathcal{P}}$ and $\vec{\mathcal{P}}'$ respectively, and the new (partially filled) sketch is now $\mathcal{S}_f[\vec{\mathcal{P}}/\vec{\mathcal{H}}, \vec{\mathcal{P}}'/\vec{\mathcal{H}}']$. Since there is an opportunity for rejecting this partially filled sketch, we again check whether $\mathcal{S}_f[\vec{\mathcal{P}}/\vec{\mathcal{H}}, \vec{\mathcal{P}}'/\vec{\mathcal{H}}']$ is consistent with the input-output examples using deduction.

Example 12. Consider hypothesis $\mathcal{H}$ from Figure 6, the input table $\mathbf{T}_1$ from Figure 8, and the output table $\mathbf{T}_3$ from Figure 15. We can successfully convert this hypothesis into the sketch $\lambda x. ?_0^x(?_1^x(?_3^x @ (x, \mathbf{T}_1), ?_4), ?_2)$. Since FILLSKETCH is bottom-up, it first tries to fill hole $?_4$. In this case, suppose that we try to instantiate hole $?_4$ with the predicate $\text{age} > 12$ using rule (1) from Figure 14. However, when we call DEDUCE on the partially-completed sketch

$\lambda x. ?_0^\pi (?_1^\sigma (?_3 @ (x, T_1), \text{age} > 12), ?_2), ?_1$ is refined as T_4 in Figure 15 and we obtain the following constraint:

$$\begin{aligned} & ?_1.\text{row} < ?_3.\text{row} \wedge ?_1.\text{col} = ?_3.\text{col} \wedge ?_0.\text{row} = ?_1.\text{row} \wedge \\ & ?_0.\text{col} < ?_1.\text{col} \wedge x_1 = ?_3 \wedge x_1.\text{row} = 3 \wedge x_1.\text{col} = 4 \wedge \\ & y = ?_0 \wedge y.\text{row} = 2 \wedge y.\text{col} = 3 \wedge \underline{?_1.\text{col} = 4} \wedge \underline{?_1.\text{row} = 1} \end{aligned}$$

Note that the last two conjuncts (underlined) are obtained using partial evaluation. Since this formula is unsatisfiable, we can reject this hypothesis without having to fill hole $?_2$.

8. Implementation

We have implemented our synthesis algorithm in a tool called MORPHEUS, written in C++. MORPHEUS uses the Z3 SMT solver [8] with the theory of Linear Integer Arithmetic for checking the satisfiability of constraints generated by our deduction engine.

Recall from Section 5 that MORPHEUS uses a cost model for picking the “best” hypothesis from the worklist. Inspired by previous work on code completion [28], we use a cost model based on a statistical analysis of existing code. Specifically, MORPHEUS analyzes existing code snippets that use components from Λ_T and represents each snippet as a ‘sentence’ where ‘words’ correspond to components in Λ_T . Given this representation, MORPHEUS uses the 2-gram model in SRILM [34] to assign a score to each hypothesis. Specifically, we train our language model by collecting approximately 15,000 code snippets from Stackoverflow using the search keywords `tidyr` and `dplyr`. For each code snippet, we ignore its control flow and represent it using a “sentence” where each “word” corresponds to an API call. Based on this training data, the hypotheses in the worklist W from Algorithm 1 are then ordered using the scores obtained from the n -gram model.

Following the *Occam’s razor* principle, MORPHEUS explores hypotheses in increasing order of size. However, if the size of the correct hypothesis is a large number k , MORPHEUS may end up exploring many programs before reaching length k . In practice, we have found that a better strategy is to exploit the inherent parallelism of our algorithm. Specifically, MORPHEUS uses multiple threads to search for solutions of different sizes and terminates as soon as any thread finds a correct solution.

9. Evaluation

To evaluate our method, we collected 80 data preparation tasks, all of which are drawn from discussions among R users on Stackoverflow. The MORPHEUS project webpage [4] contains (i) the Stackoverflow post for each benchmark, (ii) an input-output example, and (iii) the solution synthesized by MORPHEUS.

Our evaluation aims to answer the following questions:

Q1. Can MORPHEUS successfully automate real-world data preparation tasks and what is its running time?

Q2. How big are the benefits of SMT-based deduction and partial evaluation in the performance of MORPHEUS?

Q3. How complex are the data preparation tasks that can be successfully automated using MORPHEUS?

Q4. Are there existing synthesis tools that can also automate the data preparation tasks supported by MORPHEUS?

To answer these questions, we performed a series of experiments on the 80 data preparation benchmarks, using the input-output examples provided by the authors of the Stackoverflow posts. In these experiments, we use the table transformation components provided by two popular table manipulation libraries, namely `tidyr` and `dplyr`. The value transformers we use in our evaluation include standard comparison operators such as `<`, `>` as well as aggregate functions like `MEAN` and `SUM`. In total, our experiments make use of a total of 20 different components. All experiments are conducted on an Intel Xeon(R) computer with an E5-2640 v3 CPU and 32G of memory, running the Ubuntu 16.04 operating system and using a timeout of 5 minutes.

Summary of results. The results of our evaluation are summarized in Figure 16. Here, the “*Description*” column provides a brief English description of each category, and the column “#” shows the number of benchmarks in each category. The “*No deduction*” column indicates the running time of a version of MORPHEUS that uses purely enumerative search without deduction. (This basic version still uses the statistical analysis described in Section 8 to choose the “best” hypothesis.) The columns labeled “*Spec 1*” and “*Spec 2*” show variants of MORPHEUS using two different component specifications. Specifically, *Spec 1* is less precise and only constrains the relationship between the number of rows and columns, as shown in Table 1. On the other hand, *Spec 2* is strictly more precise than *Spec 1* and also uses other information, such as cardinality and number of groups.

Performance. As shown in Figure 16, the full-fledged version of MORPHEUS (using the more precise component specifications) can successfully synthesize 78 out of the 80 benchmarks and times out on only 2 problems. Hence, overall, MORPHEUS achieves a success rate of 97.5% within a 5-minute time limit. MORPHEUS’s median running time on these benchmarks is 3.59 seconds, and 86.3% of the benchmarks can be synthesized within 60 seconds. However, it is worth noting that running time is actually dominated by the R interpreter: MORPHEUS spends roughly 68% of the time in the R interpreter, while using only 15% of its running time to perform deduction (i.e., solve SMT formulas). Since the overhead of the R interpreter can be significantly reduced with sufficient engineering effort, we believe there is considerable room for improving MORPHEUS’s running time. However, even in its current form, these results show that MORPHEUS is practical enough to automate a diverse class of data preparation tasks within a reasonable time limit.

Category	Description	#	No deduction		Spec 1		Spec 2	
			#Solved	Time	#Solved	Time	#Solved	Time
C1	Reshaping dataframes from either “long” to “wide” or “wide” to “long”	4	2	198.14	4	15.48	4	6.70
C2	Arithmetic computations that produce values not present in the input tables	7	6	5.32	7	1.95	7	0.59
C3	Combination of <i>reshaping</i> and <i>string manipulation</i> of cell contents	34	28	51.01	31	6.53	34	1.63
C4	Reshaping and arithmetic computations	14	9	162.02	10	90.33	12	15.35
C5	Combination of <i>arithmetic computations</i> and <i>consolidation</i> of information from multiple tables into a single table	11	7	8.72	10	3.16	11	3.17
C6	Arithmetic computations and string manipulation tasks	2	1	280.61	2	49.33	2	3.03
C7	Reshaping and consolidation tasks	1	0	✗	1	135.32	1	130.92
C8	Combination of <i>reshaping</i> , <i>arithmetic computations</i> and <i>string manipulation</i>	6	1	✗	3	198.42	6	38.42
C9	Combination of <i>reshaping</i> , <i>arithmetic computations</i> and <i>consolidation</i>	1	0	✗	0	✗	1	97.3
Total		80	54 (67.5%)	95.53	68 (85.0%)	8.57	78 (97.5%)	3.59

Figure 16. Summary of experimental results. All times are median in seconds and ✗ indicates a timeout (> 5 minutes).

Figure 17. Cumulative running time of MORPHEUS

Impact of deduction. As Figure 16 shows, deduction has a huge positive impact on the algorithm. The basic version of MORPHEUS that does not perform deduction times out on 32.5% of the benchmarks and achieves a median running time of 95.53 seconds. On the other hand, if we use the coarse specifications given by *Spec 1*, we already observe a significant improvement. Specifically, using *Spec 1*, MORPHEUS can successfully solve 68 out of the 80 benchmarks, with a median running time of 8.57 seconds. These results show that even coarse and easy-to-write specifications can have a significant positive impact on synthesis.

Impact of partial evaluation. Figure 17 shows the cumulative running time of MORPHEUS with and without partial evaluation. Partial evaluation significantly improves the performance of MORPHEUS, both in terms of running time and the number of benchmarks solved. In particular, without par-

Figure 18. Impact of language model

tial evaluation, MORPHEUS can only solve 62 benchmarks with median running time of 34.75 seconds using *Spec 1* and 64 benchmarks with median running time of 17.07 seconds using *Spec 2*. When using partial evaluation, MORPHEUS can prune 72% of the partial programs without having to fill all holes in the sketch, thereby resulting in significant performance improvement.

Impact of language model. As described in Section 8, MORPHEUS uses a statistical language model (namely 2-grams) for choosing the most promising hypothesis in its worklist. Even though the idea of using statistical language models is not a contribution of this paper and is inspired by the prior work of Raychev et al. [28], we nevertheless evaluate its impact on our benchmark set consisting of various data preparation tasks. Specifically, Figure 18 shows the percentage of benchmarks solved by MORPHEUS with and without a language model for ordering the hypotheses. As shown in Figure 18, the use of the language model has a significant positive impact on the performance of MORPHEUS. Specifically, while MORPHEUS can solve 97.5% of

Figure 19. Comparison with SQLSYNTHESIZER

the benchmarks using the statistical language model, it is only able to solve 76.25% of the benchmarks without the 2-gram model. However, it is worth noting that the statistical language model alone is not sufficient for solving many of our benchmarks. In particular, if we disable the deductive reasoning capabilities of MORPHEUS, we can only solve 67.5% of the benchmarks. Furthermore, MORPHEUS can only solve 28.75% of the benchmarks if we disable both deduction as well as the statistical language model.

Complexity of benchmarks. To evaluate the complexity of tasks that MORPHEUS can handle, we conducted a small user study involving 9 participants. Of the participants, four are senior software engineers at a leading data analytics company and do data preparation “for a living”. The remaining 5 participants are proficient R programmers at a university and specialize in statistics, business analytics, and machine learning. We chose 5 representative examples from our 80 benchmarks and asked the participants to solve as many of them as possible within one hour. These benchmarks belong to four categories (C2, C3, C4, C7) and take between 0.22 and 204.83 seconds to be solved by MORPHEUS.

In our user study, the average participant completed 3 tasks within the one-hour time limit; however, only 2 of these tasks were solved *correctly* on average. These results suggest that our benchmarks are challenging even for proficient R programmers and expert data analysts.

Comparison with λ^2 . To demonstrate the advantages of our proposed approach over previous component-based synthesis techniques, we compared MORPHEUS with λ^2 [10], which is a general-purpose tool for synthesizing higher-order functional programs over data structures.

Since λ^2 does not have built-in support for tables, we evaluated λ^2 on the benchmarks from Figure 16 by representing each table as a list of lists. Even though we confirmed that λ^2 can synthesize very simple table transformations involve projection and selection, it was not able to successfully synthesize *any* of the benchmarks used in our evaluation. Upon further inspection, we believe that λ^2 fails to synthesize many of our benchmarks for two reasons: First, hypotheses in λ^2 are restricted to be of the form $\lambda x. F e x$, where F is a higher-order combinator, e is an expression and x is the input. However, many of our benchmarks re-

quire more general hypotheses of the form $\lambda x. F e_1 e_2$ where e_1, e_2 are arbitrary expressions. Furthermore, λ^2 can only perform deduction for a built-in set of higher-order combinators for which it is possible to infer concrete input-output examples for the sub-components. However, many of the benchmarks used in our evaluation are difficult to express concisely using the set of combinators supported by λ^2 .

Comparison with SQLSynthesizer. Since MORPHEUS is a general tool that can be used to synthesize many kinds of table transformations, we also compare it against SQLSYNTHESIZER, which is a specialized tool for synthesizing SQL queries from examples [37]. To compare MORPHEUS with SQLSYNTHESIZER, we used two different sets of benchmarks. First, we evaluated SQLSYNTHESIZER on the 80 data preparation benchmarks from Figure 16. Note that some of the data preparation tasks used in our evaluation cannot be expressed using SQL, and therefore fall beyond the scope of a tool like SQLSYNTHESIZER. Among our 80 benchmarks, SQLSYNTHESIZER was only able to successfully solve *one*.

To understand how MORPHEUS compares with SQLSYNTHESIZER on a narrower set of table transformation tasks, we also evaluated both tools on the 28 benchmarks used in evaluating SQLSYNTHESIZER [37]. To solve these benchmarks using MORPHEUS, we used the same input-output tables as SQLSYNTHESIZER and used a total of eight higher-order components that are relevant to SQL. As shown in Figure 19, MORPHEUS also outperforms SQLSYNTHESIZER on these benchmarks. In particular, MORPHEUS can solve 96.4% of the SQL benchmarks with a median running time of 1 second whereas SQLSYNTHESIZER can solve only 71.4% with a median running time of 11 seconds.

10. Related Work

In this section, we compare and contrast our approach with prior synthesis techniques.

PBE for table transformations. This paper is related to a line of work on programming-by-example (PBE) [5, 6, 10, 12, 17, 21, 22, 24, 25, 27, 36]. Of particular relevance are PBE techniques that focus on table transformations [6, 17, 22, 37]. Among these techniques, FLASHEXTRACT and FLASHRELATE address the specific problem of extracting structured data from spreadsheets and do not consider a general class of table transformations. More closely related are Harris and Gulwani’s work on synthesis of spreadsheet transformations [17] and Zhang et al.’s work on synthesizing SQL queries [37]. Our approach is more general than these methods in that they use DSLs with a fixed set of primitive operations (components), whereas our approach takes a set of components as a *parameter*. For instance, Zhang et al. cannot synthesize programs that perform table reshaping while Harris et al. supports data reshaping, but not computation or consolidation. Hence, these approaches cannot automate many of the data preparation tasks that we consider.

Data wrangling. Another term for data preparation is “*data wrangling*”, and prior work has considered methods to facilitate such tasks. For instance, WRANGLER is an interactive visual system that aims to simplify data wrangling [15, 20]. OPENREFINE is a general framework that helps users perform data transformations and clean messy data. Tools such as WRANGLER and OPENREFINE facilitate a larger class of data wrangling tasks than MORPHEUS, but they do not automatically synthesize table transformations from examples.

Synthesis using deduction and search. Our work builds on recent synthesis techniques that combine enumeration and deduction [5, 10, 22, 24, 36]. The closest work in this space is λ^2 , which synthesizes functional programs using deduction and cost-directed enumeration [10]. Like λ^2 , we differentiate between higher-order and first-order combinators and use deduction to prune partial programs. However, the key difference from prior techniques is that our deduction capabilities are not customized to a specific set of components. For example, λ^2 only supports a fixed set of higher-order combinators and uses “baked-in” deductive reasoning to reject partial programs. Furthermore, as described in Section 9, λ^2 restricts hypotheses to be of a certain shape and can only perform deduction in cases where it is possible to infer concrete input-output examples for the nested hypotheses. In contrast, our approach supports any higher-order component and can utilize arbitrary first-order specifications to reject hypotheses using SMT solving.

Also related is FLASHMETA, which gives a generic method for constructing example-driven synthesizers for user-defined DSLs [27]. The methodology we propose in this paper is quite different from FLASHMETA. FLASHMETA uses version space algebras to represent *all* programs consistent with the examples and employs deduction to decompose the synthesis task. In contrast, we use enumerative search to find *one* program that satisfies the examples and use SMT-based deduction to reject partial programs.

Component-based synthesis. Component-based synthesis refers to generating (straight-line) programs from a set of components, such as methods provided by an API [9, 14, 18, 19, 23, 30]. Some of these efforts [14, 18] use an SMT-solver to *search* for a composition of components. In contrast, our approach uses an SMT-solver as a *pruning tool* in enumerative search and does not require precise specifications of components. Another related work in this space is SYPET [9], which searches for well-typed programs using a Petri net representation. Similar to this work, SYPET can also work with any set of components and decomposes synthesis into two separate sketch generation and sketch completion phases. However, both the application domains (Java APIs vs. table transformations) and the underlying techniques (Petri net reachability vs. SMT-based deduction) are very different. Finally, another related approach is Big λ [30], which can synthesize non-trivial data-parallel programs using a set of pre-defined components. However,

unlike MORPHEUS, Big λ does not incorporate deductive reasoning to prune the search space.

Synthesis as type inhabitation. Our approach also resembles prior work that has framed synthesis as type inhabitation [11, 16, 24, 26]. Of these approaches, INSYNTH [16] is type-directed rather than example-directed. MYTH [24] and its successors [11] cast type- and example-directed synthesis as type inhabitation in a refinement type system. In contrast to these techniques, our approach only enumerates type inhabitants in the context of sketch completion and uses table contents to finitize the universe of constants.

Another work that is closely related to MORPHEUS is SYNQUID [26], which takes advantage of recent advances in polymorphic refinement types [29, 35]. Similar to our approach, SYNQUID also adopts a type-directed SMT-based deduction system to prune its search space. However, unlike our system which can work with any incomplete (over-approximate) specification, SYNQUID requires precise specifications of the underlying components. In other words, SYNQUID fails to synthesize the desired program if the component specifications are over-approximate. Since it is difficult to write precise specifications of many library methods, we believe that MORPHEUS’s ability to perform lightweight deduction using incomplete specifications can be useful in many different contexts.

Sketch. In *program sketching*, the user provides a partial program containing holes, which are completed by the synthesizer in a way that respects user-provided invariants (e.g., assertions) [31–33]. While we also use the term “*sketch*” to denote partial programs with unknown expressions, the holes in our program sketches can be arbitrary expressions over first-order components. In contrast, holes in the SKETCH system typically correspond to constants [33]. Furthermore, our approach automatically generates program sketches rather than requiring the user to provide the sketch.

11. Conclusion

We have presented a new component-based synthesis algorithm that combines type-directed enumerative search with lightweight SMT-based deduction and partial evaluation. Given a set of components equipped with over-approximate logical specifications, our approach automatically infers logical specifications of partial programs and uses SMT-based reasoning to prune the search space. Our approach further increases the power of its deductive reasoning engine by employing partial evaluation. We have applied the proposed ideas to automate a large class of data preparation tasks that involve table consolidation and reshaping. As shown in our experimental evaluation, our tool, MORPHEUS, can automate challenging data wrangling tasks that are difficult even for proficient R programmers. Our tool is publicly available [4] and will also be released as an RStudio plug-in.

References

- [1] Motivating Example 1. <http://stackoverflow.com/questions/30399516/complex-data-reshaping-in-r>. Accessed 27-Mar-2017.
- [2] Motivating Example 2. <http://stackoverflow.com/questions/33207263/finding-proportions-in-flights-dataset-in-r>. Accessed 27-Mar-2017.
- [3] Motivating Example 3. <http://stackoverflow.com/questions/32875699/how-to-combine-two-data-frames-in-r-see-details>. Accessed 27-Mar-2017.
- [4] Morpheus. <https://utopia-group.github.io/morpheus/>. Accessed 27-Mar-2017.
- [5] A. Albarghouthi, S. Gulwani, and Z. Kincaid. Recursive Program Synthesis. In *Proc. International Conference on Computer Aided Verification*, pages 934–950. Springer, 2013.
- [6] D. W. Barowy, S. Gulwani, T. Hart, and B. G. Zorn. FlashRelate: extracting relational data from semi-structured spreadsheets using examples. In *Proc. Conference on Programming Language Design and Implementation*, pages 218–228. ACM, 2015.
- [7] T. Dasu and T. Johnson. *Exploratory data mining and data cleaning*, volume 479. John Wiley & Sons, 2003.
- [8] L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *Proc. Tools and Algorithms for Construction and Analysis of Systems*, pages 337–340. Springer, 2008.
- [9] Y. Feng, R. Martins, Y. Wang, I. Dillig, and T. Reps. Component-Based Synthesis for Complex APIs. In *Proc. Symposium on Principles of Programming Languages*. ACM, 2017.
- [10] J. K. Feser, S. Chaudhuri, and I. Dillig. Synthesizing data structure transformations from input-output examples. In *Proc. Conference on Programming Language Design and Implementation*, pages 229–239. ACM, 2015.
- [11] J. Frankle, P. Osera, D. Walker, and S. Zdancewic. Example-directed synthesis: a type-theoretic interpretation. In *Proc. Symposium on Principles of Programming Languages*, pages 802–815. ACM, 2016.
- [12] S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *Proc. Symposium on Principles of Programming Languages*, pages 317–330. ACM, 2011.
- [13] S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, pages 317–330. ACM, 2011.
- [14] S. Gulwani, S. Jha, A. Tiwari, and R. Venkatesan. Synthesis of loop-free programs. In *Proc. Conference on Programming Language Design and Implementation*, pages 62–73. ACM, 2011.
- [15] P. J. Guo, S. Kandel, J. M. Hellerstein, and J. Heer. Proactive Wrangling: Mixed-initiative End-user Programming of Data Transformation Scripts. In *Proc. Symposium on User Interface Software and Technology*, pages 65–74. ACM, 2011.
- [16] T. Gvero, V. Kuncak, I. Kuraj, and R. Piskac. Complete completion using types and weights. In *Proc. Conference on Programming Language Design and Implementation*, pages 27–38. ACM, 2013.
- [17] W. R. Harris and S. Gulwani. Spreadsheet table transformations from examples. In *Proc. Conference on Programming Language Design and Implementation*, pages 317–328. ACM, 2011.
- [18] S. Jha, S. Gulwani, S. Seshia, and A. Tiwari. Oracle-guided component-based program synthesis. In *Proc. International Conference on Software Engineering*, pages 215–224. IEEE, 2010.
- [19] T. A. Johnson and R. Eigenmann. Context-sensitive domain-independent algorithm composition and selection. In *Proc. Conference on Programming Language Design and Implementation*, pages 181–192. ACM, 2006.
- [20] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer. Wrangler: Interactive visual specification of data transformation scripts. In *Proc. International Conference on Human Factors in Computing Systems*, pages 3363–3372. ACM, 2011.
- [21] E. Kitzelmann. A combined analytical and search-based approach for the inductive synthesis of functional programs. *Künstliche Intelligenz*, 25(2):179–182, 2011.
- [22] V. Le and S. Gulwani. FlashExtract: a framework for data extraction by examples. In *Proc. Conference on Programming Language Design and Implementation*, pages 542–553. ACM, 2014.
- [23] D. Mandelin, L. Xu, R. Bodík, and D. Kimelman. Jungloid mining: helping to navigate the API jungle. In *Proc. Conference on Programming Language Design and Implementation*, pages 48–61. ACM, 2005.
- [24] P.-M. Osera and S. Zdancewic. Type-and-example-directed program synthesis. In *Proc. Conference on Programming Language Design and Implementation*, pages 619–630. ACM, 2015.
- [25] D. Perelman, S. Gulwani, D. Grossman, and P. Provost. Test-driven synthesis. In *Proc. Conference on Programming Language Design and Implementation*, page 43. ACM, 2014.
- [26] N. Polikarpova, I. Kuraj, and A. Solar-Lezama. Program synthesis from polymorphic refinement types. In *Proc. Conference on Programming Language Design and Implementation*, pages 522–538. ACM, 2016.
- [27] O. Polozov and S. Gulwani. FlashMeta: A framework for inductive program synthesis. In *Proc. International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 107–126. ACM, 2015.
- [28] V. Raychev, M. Vechev, and E. Yahav. Code completion with statistical language models. In *Proc. Conference on Programming Language Design and Implementation*, pages 419–428. ACM, 2014.
- [29] P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *Proc. Conference on Programming Language Design and Implementation*, pages 159–169. ACM, 2008.
- [30] C. Smith and A. Albarghouthi. Mapreduce program synthesis. In *Proc. Conference on Programming Language Design and Implementation*, pages 326–340. ACM, 2016.

- [31] A. Solar-Lezama, R. M. Rabbah, R. Bodík, and K. Ebcioglu. Programming by sketching for bit-streaming programs. In *Proc. Conference on Programming Language Design and Implementation*, pages 281–294. ACM, 2005.
- [32] A. Solar-Lezama, L. Tancau, R. Bodík, S. Seshia, and V. Saraswat. Combinatorial sketching for finite programs. In *Proc. International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 404–415. ACM, 2006.
- [33] A. Solar-Lezama, G. Arnold, L. Tancau, R. Bodík, V. A. Saraswat, and S. A. Seshia. Sketching stencils. In *Proc. Conference on Programming Language Design and Implementation*, pages 167–178. ACM, 2007.
- [34] A. Stolcke. SRILM - an extensible language modeling toolkit. In *Proc. International Conference on Spoken Language Processing*, pages 901–904. ISCA, 2002.
- [35] P. Vekris, B. Cosman, and R. Jhala. Refinement types for typescript. In *Proc. Conference on Programming Language Design and Implementation*, pages 310–325. ACM, 2016.
- [36] N. Yaghmazadeh, C. Klinger, I. Dillig, and S. Chaudhuri. Synthesizing transformations on hierarchically structured data. In *Proc. Conference on Programming Language Design and Implementation*, pages 508–521. ACM, 2016.
- [37] S. Zhang and Y. Sun. Automatically synthesizing sql queries from input-output examples. In *Proc. International Conference on Automated Software Engineering*, pages 224–234. IEEE, 2013.