



# Refinement Types for Visualization

Jingtao Xia\*  
jingtaoxia@cs.ucsb.edu  
University of California, Santa  
Barbara  
USA

Junrui Liu\*  
junrui@cs.ucsb.edu  
University of California, Santa  
Barbara  
USA

Nicholas Brown  
nobrown@sbcglobal.net  
University of California, Santa  
Barbara  
USA

Yanju Chen  
yanju@cs.ucsb.edu  
University of California, Santa  
Barbara  
USA

Yu Feng  
yufeng@cs.ucsb.edu  
University of California, Santa  
Barbara  
USA

## Abstract

Visualizations have become crucial in the contemporary data-driven world as they aid in exploring, verifying, and sharing insights obtained from data. In this paper, we propose a new paradigm of visualization synthesis based on *refinement types*. Besides input-output examples, users can optionally use refinement-type annotations to constrain the range of valid values in the example visualization or to express complex interactions between different visual components. Our system's outputs include both data transformation and visualization programs that are consistent with refinement-type specifications. To mitigate the scalability challenge during the synthesis process, we introduce a new visualization synthesis algorithm that uses lightweight bidirectional type checking to prune the search space. As we demonstrate experimentally, this new synthesis algorithm results in significant speed-up compared to prior work.

We have implemented the proposed approach in a tool called CALICO and evaluated it on 40 visualization tasks collected from online forums and tutorials. Our experiments show that CALICO can solve 98% of these benchmarks and, among those benchmarks that can be solved, the desired visualization is among the top-1 output generated by CALICO. Furthermore, CALICO takes an average of 1.56 seconds to generate the visualization, which is 50 times faster than VISER, a state-of-the-art synthesizer for data visualization.

## ACM Reference Format:

Jingtao Xia, Junrui Liu, Nicholas Brown, Yanju Chen, and Yu Feng. 2024. Refinement Types for Visualization. In *39th IEEE/ACM International Conference on Automated Software Engineering (ASE '24)*, October 27–November 1, 2024, Sacramento, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3691620.3695550>

\*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.  
*ASE '24, October 27–November 1, 2024, Sacramento, CA, USA*  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1248-7/24/10  
<https://doi.org/10.1145/3691620.3695550>

## 1 Introduction

Visualizations have become crucial in the contemporary data-driven world as they aid in exploring, verifying, and sharing insights obtained from data. With the widespread adoption of challenging visualization tasks across various application domains, there has been a surge in the development of multiple libraries that strive to simplify intricate visualization tasks. For example, in the past few years, dozens of visualization libraries have emerged in popular programming languages such as R, Python, and JavaScript. Moreover, there has been a flurry of research focused on designing programming systems such as D3 [1] and Vega-Lite [23], aimed at enhancing real-world visualization tasks.

To help data scientists visualize raw data that is exploding in quantity, there has been a growing interest in using program synthesis to automatically generate visualization programs from user demonstrations. One flavor of such demonstrations is input-output (IO) examples [29]: the user provides tabular input data and demonstrates how to visualize a small number of data points. However, for many visualizations that require complex computations, concrete examples are often insufficient for fully expressing user intent, leading to *overfitting*. Moreover, existing synthesizers (e.g., Viser [28]) based on input-output examples require the user to provide the exact values (e.g. height of a bar, or x-coordinate of a data point) via laborious manual calculation, which hinders the adoption of automated synthesizers.

In this paper, we propose a new paradigm of visualization synthesis based on *refinement types*. Refinement types are types endowed with logical formulae that constrain values; for example,  $\{v : \text{Int} \mid 0 < v\}$  stands for positive integers. Besides IO examples, the user of our system can also use refinement-type annotations to constrain the range of valid values in the example visualization or to express complex interactions between different visual components. Our system's outputs include both data transformation and visualization programs that are consistent with refinement-type specifications.

While there has been recent work on automating visualization tasks by reducing them to programming-by-example [28, 29] for table transformations, these techniques focus on cases where the specification is a pair of input and output tables. In contrast, the refinement type specification in our setting could be a partial output table whose elements (e.g., cells, column names, etc.) are refined with *logical qualifiers*  $\phi$ . Therefore, pruning strategies used

in prior work are not effective in our setting due to the lack of fine-grained handling of refinement-type annotations. However, a general refinement-type-based synthesizer [20] will not work for our case because it requires users to provide a *precise* semantic specification for each construct in the domain-specific language (DSL) for table/visualization transformation. To deal with this challenge, our key insight is to refine each type of construct in our visualization DSL with logical formulas that *over-approximate* the computational constraints on both tables and arguments, including those on how the output columns and rows are produced. These logical formulas are framework-independent, as they are formulated from the mathematical definitions of operators. Then we introduce a new visualization synthesis algorithm that uses lightweight bidirectional type checking to prune the search space. As we demonstrate experimentally, this new algorithm results in much faster synthesis performance compared to prior work [28, 29] for automating visualization tasks. We have implemented the proposed approach in a new tool called CALICO and evaluated it on 40 visualization tasks collected from online forums and tutorials. Our experiments show that CALICO can solve 98% of these benchmarks and, among those benchmarks that can be solved, the desired visualization is the top-1 output generated by CALICO. Furthermore, CALICO takes an average of 1.56 seconds to generate the visualization, which is 50 times faster than VISER, a state-of-the-art synthesizer for data visualization. We believe that CALICO is fast enough to be practically useful to prospective users.

To summarize, this paper makes the following key contributions:

- We introduce a rich and expressive specification language for data visualization based on refinement types.
- We propose a scalable algorithm for synthesizing table transformations using refinement types. Our algorithm employs lightweight bidirectional refinement type checking to effectively prune the search space.
- We evaluate our approach on over 40 tasks collected from online forums and tutorials and show that CALICO significantly outperforms prior work on data visualization.

## 2 Overview

In our CALICO framework, the user provides an input table and a demonstration of how to visualize example data points via a so-called *visual trace*. Our tool then synthesizes a program that, when evaluated on the input table, produces a visualization consistent with the user-provided visual trace. The synthesized program consists of a *table program* and a *visual program*: the former applies a sequence of table operators (e.g., projection, filtering) to transform the input table into a final table containing values needed for the visualization; the latter program renders the final table using various charts (e.g., bar charts). Since synthesis of visualization programs from a is standard [28, 29], in the remainder of the paper we focus on table transformation, although our examples may incorporate visual programs for clarity.

As an example, consider the visualization task shown in Figure 1. The user may want to visualize the input table  $t_{in}$  as the bar chart shown on the right, where each bar represents the percentage of objects with feature  $A$  for each condition. The desired program first

applies table transformation operators `spread` and `mutate` to obtain output table  $t_{out}$  that contains columns *condition* and *percentage*. Then, the visual program renders a bar chart using those columns.

How can the user demonstrate the intended visualization using visual traces? In previous works, such as Falx [29], the visual traces can only contain concrete values. That is, the user must show the exact height of the first bar with the trace  $\text{Bar}(x = 1, y = 0.667)$ . However, in order to compute the correct value of  $y$ , the user has to (i) manually locate rows in the table for which `condition = 1`, (ii) extract the count value for each feature, and (iii) use a calculator to perform the required arithmetic, a non-trivial task. As a result, the user would have manually performed the complex `spread` and `mutate` operations that would appear in the desired program. This process can quickly become unmanageable for larger tables and more complex visualizations.

A unique aspect of CALICO is that the arduous and error-prone demonstration of concrete values is replaced by more intuitive constraints in the visual traces, encoded as *refinement types*. In this example, the user can simply specify that  $x$  comes from column *condition* and  $y$  is a percentage between 0 and 1, using the visual trace  $\text{Bar}(x \in \text{condition}, 0 < \text{percentage} < 1)$ . Our tool automatically transforms this trace into a refinement type on the final table:

$$\begin{aligned} \mathcal{T}_{out} : \langle x :: \{v : \text{Int} \mid v < \{\text{condition}\}\}, \\ y :: \{v : \text{Real} \mid 0 < v \wedge v < 1\}, \rangle \end{aligned}$$

where  $v < \{\text{condition}\}$  indicates a data-flow lineage from the input column *condition* to the output column  $x$ . Here, the type  $\mathcal{T}_{out}$  describes an output table with at least two columns called  $x$  and  $y$ , each described by a refinement type. For example, column  $y$  has a refinement type whose base type is `Real`, and refined by the predicate  $P(v) \equiv 0 < v \wedge v < 1$ .

### 2.1 Synthesis via Bidirectional Type Inference

Given the input table and the type  $\mathcal{T}_{out}$  of the output table, CALICO will find a program that generates a final table of type  $\mathcal{T}_{out}$ . However, existing techniques for synthesizing programs from refinement types (e.g., SYNQUID [20]) are insufficient for solving this problem: they assume a precise semantic specification of each language construct. This assumption is known to be problematic in the table program synthesis domain [7].

Instead, we propose a novel algorithm based on the ideas of *bidirectional analysis* [19]. Our key insight is that using refinement types, despite the difficulty of accurately encoding table operations semantically using refinement types, we can define *abstract semantics* of each table operation in both forward and backward directions, even if the arguments to the operation are not fully determined. Specifically, if the input (or output) to a table operator has type  $\tau$ , we can approximate the effect of applying (or unapplying) the operator and obtain a new type  $\tau'$ ; the forward and the backward directions will eventually meet and generate a *type consistency constraint* relating  $\tau$  and  $\tau'$ . The type consistency constraint allows us to determine whether an incomplete table program is feasible, and to prune infeasible programs early to speed up the synthesis search.

Consider the following partial program as a candidate for solving the motivating example:  $t_{in} \gg \text{mutate}(\text{tmp} = \text{count} * \text{condition}) \gg$

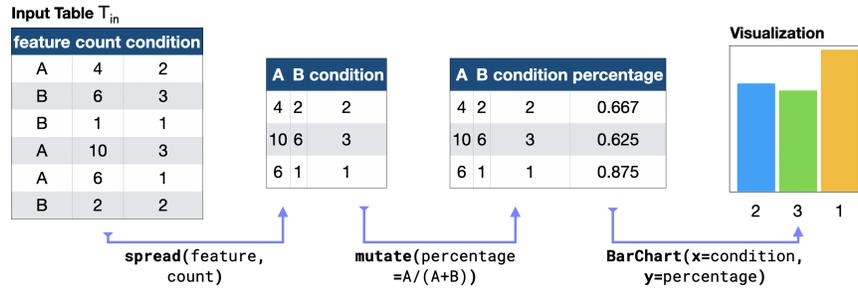


Figure 1: An example visualization task

select( $\square$ ), where  $\square$  indicates a hole yet to fill in and  $\gg$  connects operations sequentially. Although this program is incomplete, CALICO can still prove its infeasibility: regardless of what argument is supplied to select, the resulting program cannot produce the desired output table (and hence the visualization). To determine infeasibility, CALICO assigns the refinement type  $\mathcal{T}_{in}$  to the input table:

$$\mathcal{T}_{in} \equiv \langle \text{count} :: \{v : \text{Int} \mid (1 \leq v \wedge v \leq 10)\}, \\ \text{condition} :: \{v : \text{Int} \mid 1 \leq v \wedge v \leq 3\}, \dots \rangle.$$

CALICO's forward analysis infers that, after the mutate operation (which introduces a new column using the specified computation expression), the type  $\mathcal{T}'$  of the resulting table will have an additional column called tmp with the refinement type  $\{v : \text{Int} \mid 1 \leq v \wedge v \leq 30\}$ :

$$\mathcal{T}' \equiv \langle \text{tmp} :: \{v : \text{Int} \mid 1 \leq v \wedge v \leq 30\}, \\ \text{count} :: \{v : \text{Int} \mid 1 \leq v \wedge v \leq 10\}, \\ \text{condition} :: \{v : \text{Int} \mid 1 \leq v \wedge v \leq 3\}, \dots \rangle.$$

On the other hand, CALICO's backward analysis examines the operation select( $\square$ ). It sees that because the select operation can only return a table with equal or fewer columns, the type of the output table  $\mathcal{T}_{out}$  is also a valid over-approximation of the input to select. At this point, the forward and the backward analysis meet at the output of mutate and the input of select, and they infer two types that describe the same table. Thus, these two types must be consistent, and CALICO generates the following type consistency constraint:  $\mathcal{T}' \approx \mathcal{T}_{out}$ . CALICO utilizes a set of inference rules to decide whether such relations hold. In this case, because  $\mathcal{T}'$  does not have any column containing real numbers between 0 and 1 as required by  $\mathcal{T}_{out}$ , CALICO concludes that the two types are inconsistent, and the candidate program is deemed infeasible and excluded from further consideration.

### 3 Formulation

In this section, we formally define the problem of synthesizing visualization programs using refinement types. Before doing so, we will first define CALICO's table transformation language, and then present CALICO's refinement type system and its subtyping relation.

$P ::= t \gg e_1 \gg \dots \gg e_n$	Table program
$e ::=$	<b>Table operators</b>
select( $\vec{c}$ )	Projection
filter( $\sim, \vec{c}_{arg}$ )	Filtering
mutate( $c_{target}, \otimes, \vec{c}_{arg}$ )	Calculation
spread( $\vec{c}_{id}, c_{key}, c_{val}$ )	Pivoting (wider)
gather( $\vec{c}_{id}, c_{target}$ )	Pivoting (longer)
summarize( $\vec{c}_{key}, c_{target}, \ominus, \vec{c}_{arg}$ )	Summarization
$\sim ::= = \mid \neq \mid \leq \mid < \mid \bar{R}$	Predicates
$\otimes ::= + \mid - \mid * \mid / \mid \bar{M}$	Arithmetic operation
$\ominus ::= \min \mid \max \mid \text{sum} \mid \text{count} \mid \text{avg}$	Aggregate operation
$\bar{G}$	Custom operator

Figure 2: Syntax of CALICO's table transformation language

#### 3.1 Table Transformation Language

Figure 2 shows the syntax of our table transformation language, inspired by real-world languages and frameworks for data wrangling and table transformation (e.g., R and SQL). A CALICO table program  $P$  is a sequence  $e_1, \dots, e_n$  of table operators applied sequentially to an input table  $t$ . After each operation, an intermediate table is produced, and the last intermediate table is said to be the output table. Each table is a collection of ordered records (i.e., name-indexed tuples). Each tuple element is called a cell. We refer to the collection of cells indexed by the same name as a column of a table, and each item in the table collection as a row.

A table operator introduces new columns/rows to or eliminates existing columns/rows from the incoming table. The semantics of each operator is as follows:

- (1) The select operator projects the specified columns  $\vec{c}$  from the input table. That is, it retains the subset of columns named in  $\vec{c}$  and drops the remaining columns.
- (2) The filter operator retains the rows of the input table that satisfy the given predicate  $\sim$ . A predicate is a binary relation that takes two column names as inputs and includes equality, inequality, less-than-or-equal-to, less-than, or a custom predicate<sup>1</sup>:  $\bar{R}$ .

<sup>1</sup>For simplicity, we omit from our presentation non-binary relations and constants, but CALICO can be extended to support them with ease.

$\mathcal{T}$	::= $\langle \vec{\kappa}_i; \vec{\rho}_j \rangle$	Table type
$\kappa$	::= $c :: \tau$	Column type
$\rho$	::= $\langle \vec{\sigma}_i \rangle$	Row type
$\sigma$	::= $c :: \tau$	Cell type
$\tau$	::= $\{v : B \mid \phi\}$	Refinement type
$B$	::=	<b>Base types</b>
	Enum	Enumeration
	Int   Real	Numeric
$\phi$	::=	<b>Logical Qualifiers</b>
	true	Truth
	$\neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi$	Connectives
	$e <^+ \mathcal{L} \mid e <^- \mathcal{L}$	Provenance
	$e \triangleleft^+ \mathcal{O} \mid e \triangleleft^- \mathcal{O}$	Relevant operators
	$e_1 \leq e_2 \mid e_1 = e_2$	Comparison
$e$	::=	<b>Expressions</b>
	$x$	Identifier
	$c$	Constant
	$e_1 \oplus e_2$	Binary operation

Figure 3: Syntax of CALICO’s refinement type system

- (3) The spread operator is a pivot operation that takes columns  $c_{key}$  and  $c_{val}$  representing key-value pairs, and pivots the table by (i) eliminating those two columns, (ii) creating a new column named after each value in  $c_{key}$ , and (iii) filling in the cells in the newly created column using the corresponding values from  $c_{val}$ . spread effectively makes a table “wider.”
- (4) The gather operator is the inverse of spread. Given a set of columns  $\vec{c}_{target}$ , it treats the column names in  $\vec{c}_{target}$  as keys, and cells in those columns as values. It then pivots the table by eliminating  $\vec{c}_{target}$  and creating two columns that contain the keys and the values, respectively. In effect, gather makes a table “longer.”
- (5) The summarize operator first partitions the table into groups where each group contains the same  $\vec{c}_{key}$ . Next, it aggregates each group using the specified aggregate operation, which can be min, max, sum, count, avg, or a custom aggregate operation  $\bar{G}$ .

Given a table transformation program  $P$ , we say the program is a *sketch*, or a *partial program*, if some arguments to the table operators are holes ( $\square$ ) that are yet to be determined. We say that a program is *complete* if it does not contain any holes.

### 3.2 CALICO’s Refinement Type System

The purpose of CALICO’s refinement type system is twofold: (i) to present the user with an expressive language to specify the desired output; (ii) to enable the CALICO synthesizer to prune away infeasible programs from the search space as early as possible.

The language of CALICO’s refinement type is described in Figure 3:

- At the top level is the *table type*  $\mathcal{T}$ , which describes individual tables. Each table type consists of a collection of *column types*  $\kappa$  and *row types*  $\rho$ .
- A column (resp. row) type describes the properties of a column (resp. row) of the table. A column type is of form  $c :: \tau$  that maps the column name  $c$  to a refinement type  $\tau$ . A row

type  $\rho$  consists of a sequence of cell types  $\sigma$  of the form  $c :: \tau$ , where  $c$  is the column name of the cell, and  $\tau$  is the refinement type that describes the cell.

- A refinement type  $\{v : B \mid \phi\}$  consists of a base type (i.e., Enum, Int, or Real) equipped with a logical qualifier  $\phi$  that refines the set of values that can be chosen from the base type. The qualifier  $\phi$  is a formula built from true, logical connectives (i.e., negation, conjunction, and disjunction), range comparisons between expressions, and domain-specific predicates.
- CALICO’s domain-specific predicates include the *provenance* predicate  $e < \mathcal{L}$ , which tracks the lineage of data flow from the input table to the current column/cell, as well as the *related-operator* predicate  $e \triangleleft \mathcal{O}$ , which maintains the set of table operators used to compute the current column/cell. The polarity of  $<$  and  $\triangleleft$  indicates under- vs. over-approximation. That is,  $e <^+ \mathcal{L}$  means that at least the set  $\mathcal{L}$  of labels is necessary to compute  $e$ , while  $e <^- \mathcal{L}$  indicates that the maximum set of labels needed to compute  $e$  is  $\mathcal{L}$ . CALICO’s domain-specific predicates not only provide the end users with a relatively precise specification mechanism without the need for manual calculation, but also enable the synthesizer to effectively narrow down the potential candidate programs.

We use judgments of the form  $\vdash t : \mathcal{T}$  to mean table  $t$  has type  $\mathcal{T}$ . We elide the details of this judgment due to its straightforward nature.

### 3.3 Problem statement

With the definitions of the table transformation language and the refinement type system established, we can formally define our problem:

**Visualization program synthesis.** Given an input table  $t_{in}$  and a type  $\mathcal{T}_{out}$ , the *visualization program synthesis* problem is to find a table transformation program  $P$  such that, if executing  $P$  on  $t_{in}$  produces the output table  $t_{out}$ , then  $\vdash t_{out} : \mathcal{T}_{out}$ , which means that the output table satisfies the user-specified type  $\mathcal{T}_{out}$ .

## 4 Algorithm

In this section, we give an overview of our synthesis algorithm. However, because bidirectional type checking analysis is one of the main contributions of this paper, we defer a detailed discussion to Section 5.

Algorithm 1 shows CALICO’s main synthesis algorithm. The SYNTHESIZE procedure takes as parameters the input table  $t_{in}$ , the type  $\mathcal{T}_{out}$  of the output table, and an integer  $k$  that indicates the maximum length of synthesized programs. The procedure first infers the type  $\mathcal{T}_{in}$  of  $t_{in}$ ; this can be done in a straightforward way since the input table is concrete. Next, it enumerates all sketches up to length  $k$  and stores them to work list  $S$ .

The core synthesis loop (L4-15) will gradually fill out each sketch with concrete arguments until the first viable program is found. At each iteration, the search procedure analyzes the next sketch  $P^*$  in the work list.

The key insight of CALICO’s synthesis search is that even if a program is incomplete, we can *over-approximate* its behavior using types. If the over-approximated behavior does not satisfy

$$\begin{array}{c}
\frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle \quad \overline{c'_j} \subseteq \overline{c_i}}{\vdash t \gg \text{select}(\overline{c'_j}) : \langle \overline{c'_j} :: \overline{\tau_j} \rangle} \text{TF-SELECT} \quad \frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle}{\vdash t \gg \text{filter}(\_, \_) : \langle \overline{c_i} :: \overline{\tau_i} \rangle} \text{TF-FILTER} \\
\\
\frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle \quad \overline{c_i} = \overline{c_{id}} \uplus \{c_{key}\} \uplus \{c_{val}\} \quad \vDash \tau_{key} <: \{v : \text{Enum} \mid v = A_1 \vee \dots \vee v = A_l\}}{\vdash t \gg \text{spread}(\overline{c_{id}}, c_{key}, c_{val}) : \langle \overline{c_{id}} :: \overline{\tau_{id}}, A_1 :: \tau_{val}, \dots, A_l :: \tau_{val} \rangle} \text{TF-SPREAD} \\
\\
\frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle \quad \overline{c_i} = \overline{c_{id}} \uplus \overline{c_{target}} \quad \tau_{key} = \{v : \text{Enum} \mid \bigvee_{c \in \overline{c_{target}}} v = c\} \quad \vDash \tau_{target} <: \tau_{val} \text{ for all } \tau_{target} \in \overline{c_{target}} \quad c_{key}, c_{val} \text{ fresh}}{\vdash t \gg \text{gather}(\overline{c_{id}}, \overline{c_{target}}) : \langle \overline{c_{id}} :: \overline{\tau_{id}}, c_{key} :: \tau_{key}, c_{val} :: \tau_{val} \rangle} \text{TF-GATHER} \\
\\
\frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle \quad c_1, c_2 \in \overline{c_i} \quad \vDash \tau_1 <: \{v : \text{Int} \mid x \leq v\} \quad \vDash \tau_2 <: \{v : \text{Int} \mid y \leq v\} \quad c_{target} \notin \overline{c_i}}{\vdash t \gg \text{mutate}(c_{target}, +, [c_1, c_2]) : \langle \overline{c_i} :: \overline{\tau_i}, c_{target} :: \{v : \text{Int} \mid x + y \leq v\} \rangle} \text{TF-MUTATEADD} \\
\\
\frac{\vdash t : \langle \overline{c_i} :: \overline{\tau_i} \rangle \quad \overline{c_{key}} \subseteq \overline{c_i} \quad c_{target} \notin \overline{c_i}}{\vdash t \gg \text{summarize}(\overline{c_{key}}, c_{targets}, \text{count}, \emptyset) : \langle \overline{c_{key}} :: \overline{\tau_{key}}, c_{target} :: \{v : \text{Int} \mid 0 \leq v\} \rangle} \text{TF-SUMMARIZECOUNT}
\end{array}$$

Figure 4: Typing rules for forward analysis

**Algorithm 1** Main synthesis algorithm

```

1: procedure SYNTHESIZE( $t_{in}, \mathcal{T}_{out}, k$ )
2:    $\mathcal{T}_{in} \leftarrow \text{inferType}(t_{in})$ 
3:    $S \leftarrow \text{enumerateSketches}(k)$ 
4:   while notEmpty( $S$ ) do
5:      $P^* \leftarrow S.\text{pop}()$ 
6:      $\mathcal{T}_f \leftarrow \text{forward}(P^*, \mathcal{T}_{in}, t_{in})$ 
7:      $\mathcal{T}_b \leftarrow \text{backward}(P^*, \mathcal{T}_{out})$ 
8:     if  $\vDash \mathcal{T}_f \preceq \mathcal{T}_b$  then
9:       if  $P^*$  is complete and  $\vdash P^*(t_{in}) : \mathcal{T}_{out}$  then
10:        yield  $P^*$ 
11:       else
12:          $S.\text{extend}(\text{expand}(P^*))$ 
13:       end if
14:     end if
15:   end while
16:   return  $\perp$ 
17: end procedure

```

$$\begin{array}{c}
\frac{}{\vDash T <: T} \text{S-REFL} \quad \frac{\vDash T_1 <: T_2 \quad \vDash T_2 <: T_3}{\vDash T_1 <: T_3} \text{S-TRANS} \\
\\
\frac{\vdash t : T_1 \quad \vDash T_1 <: T_2}{\vdash t : T_2} \text{S-SUB} \\
\\
\frac{\forall v, [[\phi_1]] \implies [[\phi_2]] \text{ is valid}}{\vDash \{v : B \mid \phi_1\} <: \{v : B \mid \phi_2\}} \text{S-REFINE}
\end{array}$$

Figure 5: Subtyping rules

the user specification, then any completion of the program will not either. Hence, the sketch, which encodes a large amount of the search space, can be rejected early, preventing the synthesizer from further exploration and branching.

This insight materializes in CALICO's forward and backward sub-routines. The former infers the types of the intermediate tables

starting from the original input table (5.1). The latter performs inference starting from the output type and works backward (5.2). The two directions eventually meet, producing types  $\mathcal{T}_f$  and  $\mathcal{T}_b$ , respectively. Since we cannot use the table typing judgment to check for the feasibility of the current program as we have no concrete tables to work with, we instead delegate the feasibility query to a novel *table subtyping relation*, which asks whether  $\mathcal{T}_f$  is *consistent* with  $\mathcal{T}_b$  (5.3). Intuitively, this checks whether  $\mathcal{T}_f$  has at least the columns and rows required by type  $\mathcal{T}_b$ .

If the consistency relation holds, the synthesis procedure returns  $P^*$  in case it is complete (i.e. does not contain any hole) and running the program produces an output table that satisfies the output type. Otherwise, we enumerate all possible ways in which  $P^*$  can be expanded (i.e. one hole being replaced) in a breadth-first manner, and add the expanded sketches to the work list. Finally, if all programs are exhausted, the procedure reports  $\perp$  to indicate that the synthesis problem is unsatisfiable.

## 5 Pruning via Bidirectional Type Inference

As is evident from the discussion above, a key part of our synthesis algorithm is the forward and backward inference to generate type consistency constraints. These procedures are described in Figure 4 and Figure 6 using inference rules.

### 5.1 Typing Rules for Forward Analysis

The typing rules for forward analysis are shown in Figure 4. In particular, each rule computes the refinement type of the output table based on the table transformation language defined in Figure 2.

- (1) **select**: The select operator retains columns given as arguments  $\overline{c'_j}$  while removing everything else. Therefore, the refinement types of  $\overline{c'_j}$  will be the type of the output table.
- (2) **filter**: Because the filter operator does not alter columns of the input table, the refinement type of the output table should match the type of the input table.
- (3) **spread**: The spread function turns the table to a wider format. It moves cells from the  $c_{val}$  column to new columns, whose

$$\begin{array}{c}
\frac{\vdash t \gg \text{select}(\_) : \mathcal{T}}{\vdash t : \mathcal{T}} \text{ TB-SELECT} \\
\\
\frac{\vdash t \gg \text{filter}(\_, \_) : \langle \vec{c}_i :: \vec{\tau}_i \rangle \quad \tau'_i = \{v : B \mid \text{true}\} \text{ for all } i \text{ and } \tau_i = \{v : B \mid \phi\}}{\vdash t : \langle \vec{c}_i :: \vec{\tau}'_i \rangle} \text{ TB-FILTER} \\
\\
\frac{\vdash t \gg \text{spread}(\vec{c}_{id}, \_) : \langle \vec{c}_i :: \vec{\tau}_i \rangle \quad \vec{c}_i = \vec{c}_{id} \uplus \vec{c}_j \quad \vDash \tau_j <: \tau_{val} \text{ for all } j \quad \tau_{key} = \{v : \text{Enum} \mid \bigvee_{c \in \vec{c}_j} v = c\} \quad c_{key}, c_{val} \text{ fresh}}{\vdash t : \langle \vec{c}_{id} :: \vec{\tau}_{id}, c_{key} :: \tau_{key}, c_{val} :: \tau_{val} \rangle} \text{ TB-SPREAD} \\
\\
\frac{\vdash t \gg \text{gather}(\vec{c}_{id}, \_) : \langle \vec{c}_i :: \vec{\tau}_i \rangle \quad \vec{c}_i = \vec{c}_{id} \uplus \{c_{key}, c_{val}\} \quad \vDash \tau_{key} <: \{v : \text{Enum} \mid v = A_1 \vee \dots \vee v = A_l\}}{\vdash t : \langle \vec{c}_{id} :: \vec{\tau}_{id}, A_1 :: \tau_{val}, \dots, A_n :: \tau_{val} \rangle} \text{ TB-GATHER} \\
\\
\frac{\vdash t \gg \text{mutate}(c_{targets}, \_) : \langle \vec{c}_i :: \vec{\tau}_i \rangle \quad \vec{c}_i = \vec{c}_j \uplus \{c_{target}\}}{\vdash t : \langle \vec{c}_j :: \vec{\tau}_j \rangle} \text{ TB-MUTATE} \\
\\
\frac{\vdash t \gg \text{summarize}(\_, c_{targets}, \_) : \langle \vec{c}_i :: \vec{\tau}_i \rangle \quad \vec{c}_i = \vec{c}_j \uplus \{c_{target}\}}{\vdash t : \langle \vec{c}_j :: \vec{\tau}_j \rangle} \text{ TB-SUMMARIZE}
\end{array}$$

Figure 6: Typing rules for backward analysis

names come from the cells in the  $c_{key}$  column. As a result, we maintain the column refinement type for all columns in  $\vec{c}_{id}$ , which is the complement of  $\{c_{val}, c_{key}\}$ . Then, we introduce new columns for each name in  $c_{key}$ . These new columns will have the same refinement type as  $c_{val}$ , as they are subsets of  $c_{val}$ .

- (4) gather: The gather operation is the inverse of the spread operation. It transforms the table into a longer format by consolidating all values in  $\vec{c}_{target}$  into a new column,  $c_{val}$ , while the column names in  $\vec{c}_{target}$  appear correspondingly in the new column  $c_{key}$ . Therefore, the rule states that the column refinement type of  $c_{key}$  is an Enum type with elements collected from the names of  $\vec{c}_{target}$ . The column type of  $c_{val}$  is the *super-type* of all value types of cells collected from  $\vec{c}_{target}$ . (The typing rules for subtype/super-type judgment of the form  $\vDash e <: \tau$  are presented in Figure 5.) Finally, the refinement type of the output table is composed of the original columns that are not selected (i.e.,  $\vec{c}_{id}$ ) and the refinement types of new columns  $\vec{c}_{key}$  and  $\vec{c}_{val}$ .
- (5) mutate: The mutate operation generates a new column,  $c_{target}$ , whose values are obtained by performing a binary operator  $\otimes$  to the cells of two selected columns, namely  $c_1$  and  $c_2$ . Therefore, the output table's type will be the union of the type of the input table and the type of the new column  $c_{target}$ . The TF-MUTATEADD rule shows the case where the binary operator is  $+$ .
- (6) summarize: The summarize operation functions similarly to mutate, but it performs aggregate operation  $\ominus$  based on the group  $\vec{c}_{key}$ . The rule here for summarize is instantiated with the count aggregate operation. The rule states that the refinement type of the output table is comprised of the set of columns  $\vec{c}_{key}$  that is used in a group, and the  $c_{target}$  column that is aggregated upon. For the count aggregator, the rule over-approximates the value of each cell in column  $c_{target}$

by adding a logical qualifier stating that all values are non-negative.

Accordingly, these rules are applicable to general functions with an arbitrary number of arguments, providing flexibility for various use cases. The refinement type of the input table, which is the initial state of the forward analysis, can be generated easily without manual effort.

## 5.2 Typing Rules for Backward Analysis

The rules for backward analysis are presented in Figure 6. In the backward analysis, our goal is to infer the refinement types of input table  $t$  from the given table operator and the refinement type of the output table. The initial state of the backward analysis is the type of the output table specified by the user.

A naive backward analysis that explores the entire search space is unlikely to scale. To mitigate this issue, our backward typing rules work with *symbolic* program sketches whose arguments may be unknown. We will use the symbol  $\_$  to stand for the parts that have not been enumerated. However, a key design decision is that our analysis does not compute the strongest necessary preconditions to ensure that the cost of type checking does not outweigh the benefits. Intuitively, as more enumeration is performed, we can gain more information, but this may also result in exploring a larger search space.

- (1) select: This rule states that the input table has the same type as the output.
- (2) filter: To prevent the synthesizer from eagerly enumerating complex arguments of the filter operator, rule TB-FILTER over-approximates the refinement type of the input table by keeping it the same as the refinement types of the output table except for weakening the logical qualifier  $\phi$  to true.
- (3) gather: The refinement types of the incoming table for the gather operator are composed of two parts: 1) the subset of columns  $\vec{c}_{id} :: \vec{\tau}_{id}$  that is not affected by the pivoting, and 2)

$$\begin{array}{c}
\frac{\kappa'_i = \kappa_i[c \mapsto c'] \text{ for all } i \in \{1..m\} \quad \rho'_j = \rho_j[c \mapsto c'] \text{ for all } j \in \{1..n\}}{\vDash \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m\}, j \in \{1..n\}} \triangleright \langle \vec{\kappa}'_i; \vec{\rho}'_j \rangle_{i \in \{1..m\}, j \in \{1..n\}}} \text{C-RENAME} \\
\\
\frac{\vec{c}'_i, \vec{\kappa}'_i \text{ is a permutation of } \vec{c}_i, \vec{\kappa}_i \text{ by exchanging columns } k \text{ and } l \quad \rho'_j = \text{ExchangeColumn}(\rho_j, k, l) \text{ for all } j \in \{1..n\}}{\vDash \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m\}, j \in \{1..n\}} \triangleright \langle \vec{\kappa}'_i; \vec{\rho}'_j \rangle_{i \in \{1..m\}, j \in \{1..n\}}} \text{C-PERM COL} \\
\\
\frac{\vec{\rho}'_i \text{ is a permutation of } \vec{\rho}_i}{\vDash \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m\}, j \in \{1..n\}} \triangleright \langle \vec{c}_i; \vec{\rho}'_j \rangle_{i \in \{1..m\}, j \in \{1..n\}}} \text{C-PERM ROW} \\
\\
\frac{0 \leq k \quad 0 \leq l \quad \rho'_j = \text{RemoveColumn}(\rho_j, m..k) \text{ for all } j \in \{1..n\}}{\vDash \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m+k\}, j \in \{1..n+l\}} \triangleright \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m\}, j \in \{1..n\}}} \text{C-SUBTABLE} \\
\\
\frac{\vDash \kappa_i \triangleright \kappa'_i \text{ for all } i \in \{1..m\} \quad \vDash \rho_j \triangleright \rho'_j \text{ for all } j \in \{1..n\}}{\vDash \langle \vec{\kappa}_i; \vec{\rho}_j \rangle_{i \in \{1..m\}, j \in \{1..n\}} \triangleright \langle \vec{\kappa}'_i; \vec{\rho}'_j \rangle_{i \in \{1..m\}, j \in \{1..n\}}} \text{C-DEPTH} \\
\\
\frac{\vDash \tau_1 \triangleright \tau_2}{\vDash c :: \tau_1 \triangleright c :: \tau_2} \text{C-COL} \quad \frac{\vDash \sigma_1 \triangleright \sigma'_2 \text{ for all } i \in \{1..n\}}{\vDash \langle \sigma_1, \dots, \sigma_n \rangle \triangleright \langle \sigma'_1, \dots, \sigma'_n \rangle} \text{C-ROW} \quad \frac{\vDash \phi_1 \triangleright \phi_2}{\vDash \{v : B \mid \phi_1\} \triangleright \{v : B \mid \phi_2\}} \text{C-REFINE} \\
\\
\frac{[[\phi_1]] \wedge [[\phi_2]] \text{ is satisfiable}}{\vDash \phi_1 \triangleright \phi_2} \text{C-SAT} \quad \frac{\mathcal{L} \subseteq \mathcal{L}'}{\vDash v <^+ \mathcal{L} \triangleright v <^- \mathcal{L}'} \text{C-PROV} \quad \frac{\mathcal{O} \subseteq \mathcal{O}'}{\vDash (v <^+ \mathcal{O}) \triangleright (v <^- \mathcal{O}')} \text{C-OPS}
\end{array}$$

Figure 7: Rules for checking type consistency

a list of columns  $A_1, \dots, A_n$  selected by the gather operator. In particular, the column name of  $A_i$  can be derived from the enum value of the  $c_{key}$  column in the output table. The refinement type of each cell in  $A_i$  is obtained by asserting that it is the super-type of all refinement types of cells that appear in columns  $\vec{c}_j$ .

- (4) spread: The refinement types of incoming table for the spread operator contain three parts: 1) the subset of columns  $\vec{c}_{id} :: \tau_{id}$  that is not affected by the pivoting, 2) the refinement type of  $c_{key}$  column, and 3) the refinement type of  $c_{val}$  column. Since we do not know the concrete values of key and val columns, we assert that the refinement type of  $c_{key}$  column is the enum type of  $\vec{c}_j$ , i.e., all column names of the output column refinement types except for the ones appearing in  $\vec{c}_{id}$ . Finally, the refinement type of  $c_{val}$  column asserts that the refinement type of each cell in  $c_{val}$  is the supertype of all refinement types of cells that appear in columns  $\vec{c}_j$ .
- (5) mutate: In this case, we only require  $c_{target}$  to be concrete. The type of the input table for the mutate operator is obtained by removing the new column  $c_{target}$  from the type of the output table.
- (6) summarize: Similar to the mutate operation, the type of the input table for the summarize operator is over-approximated by removing the new column  $c_{target}$  from the type of the output table.

### 5.3 Type Consistency

The forward and backward analyses perform type inference starting from the both ends of a (possibly incomplete) table program  $P^*$ . The two analyses eventually meet, producing types  $\mathcal{T}_f$  and  $\mathcal{T}_b$  respectively. Those two types are used to determine the feasibility of  $P^*$  using the *type consistency* relation:

$$\vDash \mathcal{T}_f \triangleright \mathcal{T}_b,$$

which denotes that  $\mathcal{T}_f$  is *consistent with*  $\mathcal{T}_b$ . The inference rules for this relation are shown in Figure 7. Essentially,  $\mathcal{T}_1$  is consistent with  $\mathcal{T}_2$  if  $\mathcal{T}_1$  provides the columns and rows required by  $\mathcal{T}_2$ , but  $\mathcal{T}_1$  may contain additional columns or rows. In what follows, we expand upon each inference rule.

The rules C-RENAME, C-PERM COL, C-PERM ROW, and C-SUBTABLE stipulate that the consistency relation is invariant under various reshaping and alignment operations, i.e., column renaming, permutation of columns or rows, and adding columns/rows to the type on the left-hand side of  $\triangleright$ <sup>2</sup>.

Once the shapes of the two participant types has been aligned, rules C-DEPTH, C-COL, and C-ROW decompose the shape, and produce consistency obligations on refinement types to be handled by rule C-REFINE.

The C-SAT rule says that two refinement types satisfy the consistency relation if they are of the same base type, and that the *qualifier consistency* relation, denoted by  $\phi_1 \triangleright \phi_2$ , holds. In general,

<sup>2</sup>We note that C-SUBTABLE is a generalization of the standard width subtyping for record type.

rule C-SAT can be used to check qualifier consistency by encoding it into a satisfiability query of the formula  $[[\phi_1]] \wedge [[\phi_2]]$ . Any qualifier implication involving CALICO’s domain-specific qualifiers is checked using C-PROV and C-OPS rules. These two rules ensure that an under-approximation never exceeds an over-approximation.

Finally, our bidirectional type inference is sound. We use  $\vdash t \Rightarrow \mathcal{T}$  (resp.  $\vdash \mathcal{T} \Leftarrow t$ ) to denote that forward (resp. backward) inference assigns type  $\mathcal{T}$  to  $t$ .

**THEOREM 1 (SOUNDNESS).** *Let  $t$  be an arbitrary table,  $e$  be a table operator, and  $t' = [[e]](t)$ . Then,*

- (1)  $\vdash t : \mathcal{T}$  and  $\vdash t' \Rightarrow \mathcal{T}'$  implies  $\vdash t' : \mathcal{T}'$ ,
- (2)  $\vdash t' : \mathcal{T}'$  and  $\vdash \mathcal{T} \Leftarrow t$  implies  $\vdash t : \mathcal{T}$ .

## 6 Implementation

We have implemented the proposed idea in a tool called CALICO with 3,587 lines of code in Python. We use the Pandas library [18] for table transformation, and the Vega-Lite library [23] for visualization. In what follows, we elaborate on other key implementation decisions.

**Rendering visualization from tables.** Similar to Viser [28], CALICO needs to convert a table into its visualization, with additional visual properties attached, such as colors, shapes, etc. In particular, CALICO invokes the Vega-Lite [27] visualization tool to render the visualization from the resulting table. This reduces the complex visualization to a succinct format that is amenable to existing data-wrangling DSL.

**Global operator checking.** CALICO maintains a relevant set of operators in logical qualifiers for each column refinement type. However, adopting a global perspective on operator requirements can improve pruning power. Specifically, CALICO gathers the set denoted  $O_g$  of all operators mentioned in the output type  $\mathcal{T}_{out}$ . When analyzing a partial program  $P^*$ , CALICO collects all operators used in the concrete part of the program as  $O_c$ , and estimates whether the abstract part can satisfy the remaining requirements  $O_g \setminus O_c$ . This approach can be generalized into typing checking rules concerning input tables after forward analysis and initial output tables. The global operator set checking can prune some sketches without generating backward inference trees, thereby enhancing pruning efficiency.

**User-defined functions.** To enhance flexibility, CALICO allows users to define uninterpreted functions for mutate and summarize. These functions may have arbitrary annotated arity. For mutate, users need to provide a function that maps a list of values to a typed value. For summarize, users need to provide a function that maps a list of vectors (Series in pandas’ terminology) to a typed value, e.g.,  $G(x, y) = \text{sum}(x)/\text{sum}(y)$ . This distinction explains why we separate  $\overrightarrow{c_{arg}}$  from  $c_{target}$  in our table transformation language, instead of using the common aggregate function definition where the  $c_{target}$  overwrites the single provided argument.

We also introduce a generalized form of mutate called mutateG (grouped mutate), where the function can access aggregated values within a group. Unlike summarize, mutateG operates similarly to the partition operation in SQL. Users need to provide a function that maps a list of vectors to a typed vector, e.g.,  $H(x) = x/\text{sum}(x)$ . Such a grouped mutate operation will not influence subsequent operations with its group.

**Extended refinement type.** We note that CALICO’s refinement type system can be extended to handle a wider range of qualifiers. For example, the related operator qualifier  $v \triangleleft O$  denotes the *minimum* set of operators that need to be used during the computation of the current columns. However, we can easily introduce a symmetric qualifier  $v \triangleleft O_{max}$  to denote the maximum set of operators that can be used. Similar extensions could be applied to provenance, enumeration range, etc. The typing rules for bidirectional analysis can be similarly extended in a straightforward fashion.

## 7 Evaluation

In this section, we describe the results of the experimental evaluation, which is designed to answer the following research questions:

- **RQ1. Effectiveness:** Can CALICO solve more visualization tasks than state-of-the-art approaches within a given time limit?
- **RQ2. Scalability:** Does CALICO improve task-solving time compared to state-of-the-art approaches?
- **RQ3. Ablation:** How important are the individual refinement type rules for forward and backward analysis?

**Benchmarks.** We evaluate CALICO on a suite of 40 benchmarks collected from various sources, including prior work such as Viser [28] and online technical forums like StackOverflow<sup>3</sup>. The benchmark suite contains a variety of problems that require a wide range of data wrangling operations (e.g., projection, aggregation, mutation, and filtration) over the inputs, with various target visualization types (e.g., bar charts, pie charts, line charts). To ensure the quality of the collected benchmarks, we incorporate a semi-automatic semantic parsing procedure with manual checking to prove the synthesizer with a refinement annotation that is accurate and captures the precise user intent.

**Experimental Setup.** We compare CALICO with the state-of-the-art visualization synthesis tools, Viser [28], which is an example-driven synthesizer and does not support refinement-type-based reasoning.

All experiments are performed on MacBook Pro with 2.4 GHz Quad-Core Intel Core i5 processor and 16 GB of memory. The time limit for a single problem is set to 15 minutes.

### 7.1 Comparison on Effectiveness

To ensure a fair comparison between CALICO and Viser, we extend Viser to make sure that both tools: 1) take the same types of input, and 2) incorporate DSL constructs to support all benchmarks. In addition, as a baseline approach, we feed the specification of each benchmark to a large language model (LLM)<sup>4</sup> and prompt it to generate solution programs as the output. Figure 9 shows a performance comparison among CALICO, Viser, and baseline LLM in terms of the total number of benchmarks solved within a given time limit.

Across all 40 benchmarks, CALICO solves 39 (98%) of them. In comparison, Viser solves 29 (73%) of them, and is slightly better than the baseline LLM approach, which solves 27 (68%) of the benchmarks. As a result, CALICO solves 25% more benchmarks than Viser, and 30% more than LLM.

<sup>3</sup><https://stackoverflow.com/>

<sup>4</sup>We use ChatGPT 3 at the time of experiments.

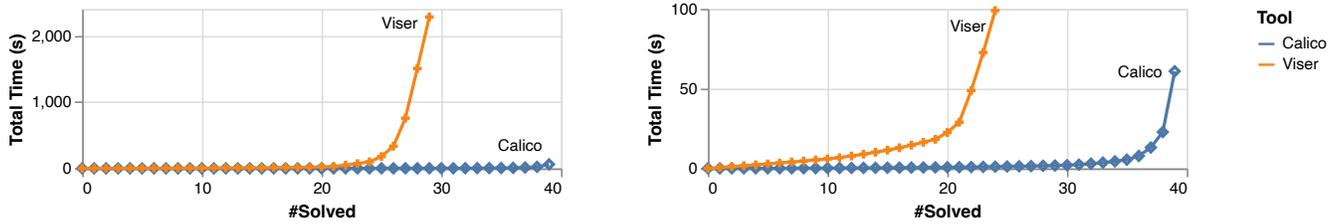


Figure 8: Scalability comparison between CALICO and VISER on visualization tasks. Each curve measures the changes of total time cost as the number of solved benchmarks grows. Left: Statistics over all benchmarks solved; Right: Statistics over all non-long-tailed benchmarks solved.

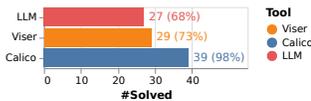


Figure 9: Effectiveness comparison between CALICO, VISER and baseline LLM on visualization tasks. Each bar is tagged with the total number of benchmarks solved as well as the percentage.

**Result for RQ1:** As is evident from these numbers, CALICO is more effective than VISER, the state-of-the-art visualization synthesis tool.

## 7.2 Comparison on Scalability

In addition to the total number of visualization tasks solved, another important evaluation metric is the time required to find the intended solution for each task. This metric is particularly important since in real-world scenarios, it reflects how much time a user has to wait before he or she is provided with a candidate to choose from.

Figure 8 shows two cactus plots comparing CALICO and VISER regarding the cumulative problem-solving time across all visualization tasks. In particular, it measures the trend over the total time required for solving each visualization task. As indicated by the plot on the left, CALICO takes *significantly* less time than VISER over all the benchmarks each tool can solve. The growth of VISER’s curve is almost exponential, CALICO remains at a low level. This shows the overall effectiveness of CALICO’s refinement type system in terms of its pruning power.

In addition to the overall comparison, we further narrow down our scope to exclude those *long-tailed* benchmarks on which both tools spend significantly more time. On the right of Figure 8 shows both trends in a more fine-grained way. We find CALICO is still showing a near-linear time growth for most of the benchmarks. Such an observation further confirms the improvement of CALICO over VISER in terms of scalability.

Our analysis on the scalability of both tools shows that the average time cost for CALICO to solve each visualization task is 1.56, while the average time cost for VISER is 78.55s. That is, CALICO is 50× faster than VISER on when it comes to visualization tasks successfully solved by themselves respectively. For benchmarks that can be solved by both tools, CALICO continues to outperform

VISER, with a 0.35s average time cost compared to VISER’s 80.41s, resulting in a 228× speed-up.

**Result for RQ2:** CALICO is scalable in that it brings a significant speed-up over the state-of-the-art tool VISER for solving each benchmark.

## 7.3 Ablation Study

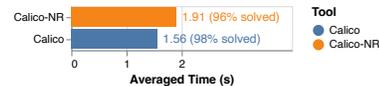


Figure 10: Performance comparison between CALICO and its ablative version CALICO-NR on visualization tasks. Each bar is tagged with the average time spent in solving and the percentage of benchmarks solved.

Recall that in Section 3.2, we introduce domain-specific predicates to prune the search space more effectively. To evaluate the effectiveness of those predicates, we perform an ablation study to compare CALICO against CALICO-NR, a Recall that a related operator keeps track of a set of table operators used to compute the target column or cell. The related operator predicate provides a relatively precise specification mechanism, which helps with pruning during synthesis.

To study the effectiveness and scalability of the related operator predicate, we measure the total number of visualization tasks solved and the average time needed to solve each of them.

Figure 10 shows the ablative results. As the full-fledged version of CALICO can solve 39 benchmarks, we can see a 3% performance drop in total number of solved benchmarks if the related operator predicate is removed (CALICO-NR), and CALICO-NR requires 22% more time on average.

**Result for RQ3:** The related operator predicate contributes to the performance of CALICO. It is effective and helpful to the overall design of the system.

## 7.4 A Discussion on Usability

We carry out a simple user study to better understand the usability of CALICO’s refinement types. In particular, participants with a basic background in data analytics are asked to write annotations for given benchmarks in addition to existing input-output examples.

We ask the participants to score the usability of the refinement type system on a scale of 1 (very easy to use) to 5 (very difficult to use). On average, the participants gave a score of 1.8, which indicates that they find refinement types annotations helpful and straightforward to provide in the majority of the cases, despite some additional learning curve.

## 7.5 Threat to Validity

There are two major threats to the validity of our conclusions, which we explain below.

*Benchmark selection.* Due to the expressiveness of the DSL, our benchmarks may not represent the actual distribution of the questions on StackOverflow. While the evaluation on the current benchmarks may not completely unveil the benefit of our approach and a representative test suite may provide a more comprehensive view, we believe our comparison and benchmarks are sufficient to show the strength of our technique. In particular, our dataset includes all difficult benchmarks from VISER [28] and additional complex benchmarks collected from StackOverflow.

*Refinement-type annotations.* Refinement types may impose an additional learning curve on the users, which could affect the usability of CALICO. Although there are cases where input-output examples are straightforward to provide, we do observe that there are many cases where refinement types are more intuitive, especially for cases involving complex arithmetic operations or relational constraints. In our experience, the refinement-type specification of all benchmarks can be easily translated from the problem description in English on StackOverflow, taking only a several minutes on average.

## 8 Related Work

There has been growing interest in automating the process of visualization generation via approaches from different research communities. In what follows, we discuss prior work in this space that is most closely related to our work.

### 8.1 Example-Driven Program Synthesis

There has been a long line of work that uses programming-by-examples (PBE) techniques to automate tedious programming tasks for various domains, e.g., string and regular expression manipulation [3, 5, 11, 33], SQL query [32, 34], table and tensor transformation [4, 6, 7], and more recently visualization synthesis VISER [28, 29, 31]. However, programming-by-example typically requires the user to encode her full intent using concrete examples, which may be either arduous or outright impossible in the presence of complex visualizations. In particular, tools like SQLSynthesizer [32] and Morpheus [7] require full input-output examples in search of a solution, and FlashFill [10] usually requires more than one pair of input-output examples to resolve ambiguity. While VISER [28] accepts partial or incomplete examples as specifications, in a real-world use case scenario, it requires the user to pick the solution that best matches her intent. Our work, CALICO, is based on example-driven program synthesis, but differs from prior work in that it allows the user to specify more details via refinement types while keeping the examples partial, which provides the flexibility for a partial but precise specification.

### 8.2 Refinement Types

Refinement types [9, 21] are type systems first introduced to enhance basic types with logical predicates. Previous works have applied variants of refinement types for program verification and program synthesis [8, 12, 13, 16, 17, 20, 26]. For example, SolType [26] builds a refinement type system for reasoning about arithmetic properties; SYNQUID builds upon bidirectional synthesis and liquid types for synthesis of recursive functions that are provably correct. Most recent work, Graphy [2] also combines refinement types with natural languages to provide finer-grained specification for visualization synthesis. Our work follows the line of work that builds upon refinement types, but focuses on the domain of visualization synthesis with a design of a scalable and effective type system.

### 8.3 Automated Visualization

Recent interest in automating visualization generation tasks has been focusing on both visualization recommendation systems and visualization exploration tools. Even though our work aligns more with the former direction, where tools like Draco [15], CompassQL [30], and ShowMe [14] prioritizes candidate visualizations according to user specifications, allowing annotation of inputs in the form of refinement types also sets our work close to the direction of visualization exploration, where tools like VisExplainer [22], Visualization-by-Sketching [24] and Polaris [25] all provides richer ways for encoding user inputs.

## 9 Conclusion

We propose a new approach to visualization synthesis based on *refinement types*. Users can specify complex interactions or calculations among visual components using refinement types. The output of our system include both data transformation and visualization programs that are consistent with the specification. To ensure scalability and usability, we introduce a new visualization synthesis algorithm that uses lightweight bidirectional type checking to prune the search space.

We have implemented the proposed approach in a new tool called CALICO and evaluated it on 40 visualization tasks collected from online forums and tutorials. Our experiments show that CALICO can solve 98% of these benchmarks and, among those benchmarks that can be solved, the desired program is among the top-1 candidate synthesized by CALICO. Furthermore, CALICO averages 1.56 seconds when solving each visualization task, which is 50 times faster than VISER, a state-of-the-art synthesizer for data visualization.

### Data Availability

An implementation of CALICO, along with the evaluation data and scripts, is available as a Docker image<sup>5</sup>.

## 10 Acknowledgement

We thank the reviewers for their helpful comments. This work is supported in part by NSF #1908494, DARPA N66001-22-2-4037, the Google Faculty Research Awards, and the Ethereum Foundation Academic Grants.

<sup>5</sup><https://hub.docker.com/r/calicosynth/calico>

## References

- [1] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D<sup>3</sup> Data-Driven Documents. *IEEE Trans. Vis. Comput. Graph.* 17, 12 (2011), 2301–2309. <https://doi.org/10.1109/TVCG.2011.185>
- [2] Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig. 2022. Type-Directed Synthesis of Visualizations from Natural Language Queries. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 144 (oct 2022), 28 pages. <https://doi.org/10.1145/3563307>
- [3] Qiaochu Chen, Xinyu Wang, Xi Ye, Greg Durrett, and Isil Dillig. 2020. Multi-Modal Synthesis of Regular Expressions. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (PLDI 2020). Association for Computing Machinery, New York, NY, USA, 487–502. <https://doi.org/10.1145/3385412.3385988>
- [4] Yanju Chen, Chenglong Wang, Osbert Bastani, Isil Dillig, and Yu Feng. 2020. Program Synthesis Using Deduction-Guided Reinforcement Learning. In *Computer Aided Verification*, Shuvendu K. Lahiri and Chao Wang (Eds.). Springer International Publishing, Cham, 587–610.
- [5] Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel rahman Mohamed, and Pushmeet Kohli. 2017. RobustFill: Neural Program Learning under Noisy I/O. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 990–998. <https://proceedings.mlr.press/v70/devlin17a.html>
- [6] Yu Feng, Ruben Martins, Osbert Bastani, and Isil Dillig. 2018. Program synthesis using conflict-driven learning. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18–22, 2018*, Jeffrey S. Foster and Dan Grossman (Eds.). ACM, 420–435.
- [7] Yu Feng, Ruben Martins, Jacob Van Geffen, Isil Dillig, and Swarat Chaudhuri. 2017. Component-Based Synthesis of Table Consolidation and Transformation Tasks from Examples. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) (PLDI 17). Association for Computing Machinery, New York, NY, USA, 422–436. <https://doi.org/10.1145/3062341.3062351>
- [8] Jonathan Frankle, Peter-Michael Osera, David Walker, and Steve Zdancewic. 2016. Example-Directed Synthesis: A Type-Theoretic Interpretation. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) (POPL '16). Association for Computing Machinery, New York, NY, USA, 802–815. <https://doi.org/10.1145/2837614.2837629>
- [9] Tim Freeman and Frank Pfenning. 1991. Refinement Types for ML. In *Proceedings of the ACM SIGPLAN 1991 Conference on Programming Language Design and Implementation* (Toronto, Ontario, Canada) (PLDI '91). Association for Computing Machinery, New York, NY, USA, 268–277. <https://doi.org/10.1145/113445.113468>
- [10] Sumit Gulwani. 2011. Automating string processing in spreadsheets using input-output examples. In *Proc. Symposium on Principles of Programming Languages*. ACM, 317–330.
- [11] Sumit Gulwani, William R. Harris, and Rishabh Singh. 2012. Spreadsheet Data Manipulation Using Examples. *Commun. ACM* 55, 8 (aug 2012), 97–105. <https://doi.org/10.1145/2240236.2240260>
- [12] Tristan Knoth, Di Wang, Nadia Polikarpova, and Jan Hoffmann. 2019. Resource-Guided Program Synthesis. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) (PLDI 2019). Association for Computing Machinery, New York, NY, USA, 253–268. <https://doi.org/10.1145/3314221.3314602>
- [13] Kenneth Knowles and Cormac Flanagan. 2009. Compositional Reasoning and Decidable Checking for Dependent Contract Types. In *Proceedings of the 3rd Workshop on Programming Languages Meets Program Verification* (Savannah, GA, USA) (PLPV '09). Association for Computing Machinery, New York, NY, USA, 27–38. <https://doi.org/10.1145/1481848.1481853>
- [14] Jock Mackinlay, Pat Hanrahan, and Chris Stolte. 2007. Show Me: Automatic Presentation for Visual Analysis. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1137–1144. <https://doi.org/10.1109/TVCG.2007.70594>
- [15] Dominik Moritz, Chenglong Wang, Greg L. Nelson, Halden Lin, Adam M. Smith, Bill Howe, and Jeffrey Heer. 2019. Formalizing Visualization Design Knowledge as Constraints: Actionable and Extensible Models in Draco. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 438–448. <https://doi.org/10.1109/TVCG.2018.2865240>
- [16] Peter-Michael Osera. 2019. Constraint-Based Type-Directed Program Synthesis. In *Proceedings of the 4th ACM SIGPLAN International Workshop on Type-Driven Development* (Berlin, Germany) (TyDe 2019). Association for Computing Machinery, New York, NY, USA, 64–76. <https://doi.org/10.1145/3331554.3342608>
- [17] Peter-Michael Osera and Steve Zdancewic. 2015. Type-and-Example-Directed Program Synthesis. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Portland, OR, USA) (PLDI '15). Association for Computing Machinery, New York, NY, USA, 619–630. <https://doi.org/10.1145/2737924.2738007>
- [18] pandas. 2023. pandas - Python Data Analysis Library. <https://pandas.pydata.org/>.
- [19] Benjamin C. Pierce and David N. Turner. 2000. Local Type Inference. *ACM Trans. Program. Lang. Syst.* 22, 1 (jan 2000), 1–44. <https://doi.org/10.1145/345099.345100>
- [20] Nadia Polikarpova, Ivan Kuraj, and Armando Solar-Lezama. 2016. Program Synthesis from Polymorphic Refinement Types. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, CA, USA) (PLDI '16). Association for Computing Machinery, New York, NY, USA, 522–538. <https://doi.org/10.1145/2908080.2908093>
- [21] Patrick Maxim Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid types. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7–13, 2008*, Rajiv Gupta and Saman P. Amarasinghe (Eds.). ACM, 159–169. <https://doi.org/10.1145/1375581.1375602>
- [22] Bahador Saket, Hannah Kim, Eli T. Brown, and Alex Endert. 2017. Visualization by Demonstration: An Interaction Paradigm for Visual Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 331–340. <https://doi.org/10.1109/TVCG.2016.2598839>
- [23] Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2017), 341–350. <https://doi.org/10.1109/TVCG.2016.2599030>
- [24] David Schroeder and Daniel F. Keefe. 2016. Visualization-by-Sketching: An Artist's Interface for Creating Multivariate Time-Varying Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 877–885. <https://doi.org/10.1109/TVCG.2015.2467153>
- [25] Chris Stolte, Diane Tang, and Pat Hanrahan. 2008. Polaris: A System for Query, Analysis, and Visualization of Multidimensional Databases. *Commun. ACM* 51, 11 (nov 2008), 75–84. <https://doi.org/10.1145/1400214.1400234>
- [26] Bryan Tan, Benjamin Mariano, Shuvendu K. Lahiri, Isil Dillig, and Yu Feng. 2022. SolType: refinement types for arithmetic overflow in solidity. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–29. <https://doi.org/10.1145/3498665>
- [27] Vega-Lite. 2019. Vega-Lite Examples. <https://vega.github.io/vega-lite/examples/>
- [28] Chenglong Wang, Yu Feng, Rastislav Bodik, Alvin Cheung, and Isil Dillig. 2019. Visualization by Example. *Proc. ACM Program. Lang.* 4, POPL, Article 49 (dec 2019), 28 pages. <https://doi.org/10.1145/3371117>
- [29] Chenglong Wang, Yu Feng, Rastislav Bodik, Isil Dillig, Alvin Cheung, and Amy J Ko. 2021. Falx: Synthesis-Powered Visualization Authoring. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 106, 15 pages. <https://doi.org/10.1145/3411764.3445249>
- [30] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2016. Towards a General-Purpose Query Language for Visualization Recommendation. In *Proceedings of the Workshop on Human-In-the-Loop Data Analytics* (San Francisco, California) (HILDA '16). Association for Computing Machinery, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/2939502.2939506>
- [31] Zhengkai Wu, Vu Le, Ashish Tiwari, Sumit Gulwani, Arjun Radhakrishna, Ivan Radiček, Gustavo Soares, Xinyu Wang, Zhenwen Li, and Tao Xie. 2022. NL2Viz: Natural Language to Visualization via Constrained Syntax-Guided Synthesis. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Singapore, Singapore) (ESEC/FSE 2022). Association for Computing Machinery, New York, NY, USA, 972–983. <https://doi.org/10.1145/3540250.3549140>
- [32] Sai Zhang and Yuyin Sun. 2013. Automatically Synthesizing SQL Queries from Input-Output Examples. In *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering* (Silicon Valley, CA, USA) (ASE '13). IEEE Press, 224–234. <https://doi.org/10.1109/ASE.2013.6693082>
- [33] Tianyi Zhang, London Lowmanstone, Xinyu Wang, and Elena L. Glassman. 2020. Interactive Program Synthesis by Augmented Examples. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 627–648. <https://doi.org/10.1145/3379337.3415900>
- [34] Xiangyu Zhou, Rastislav Bodik, Alvin Cheung, and Chenglong Wang. 2022. Synthesizing Analytical SQL Queries from Computation Demonstration. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (San Diego, CA, USA) (PLDI 2022). Association for Computing Machinery, New York, NY, USA, 168–182. <https://doi.org/10.1145/3519939.3523712>