- QIAOCHU CHEN, University of Texas at Austin, USA
- SHANKARA PAILOOR, University of Texas at Austin, USA
- ⁶ CELESTE BARNABY, University of Texas at Austin, USA
- 8 ABBY CRISWELL, University of Texas at Austin, USA
- 9 CHENGLONG WANG, Microsoft Research, USA
- 10 GREG DURRETT, University of Texas at Austin, USA
- 11 ISIL DILLIG, University of Texas at Austin, USA

12 We propose a new technique based on program synthesis for automatically generating visualizations from 13 natural language queries. Our method parses the natural language query into a refinement type specification 14 using the intents-and-slots paradigm and leverages type-directed synthesis to generate a set of visualization 15 programs that are most likely to meet the user's intent. Our refinement type system captures useful hints 16 present in the natural language query and allows the synthesis algorithm to reject visualizations that violate well-established design guidelines for the input data set. We have implemented our ideas in a tool called Graphy 17 and evaluated it on NLVCORPUS, which consists of 3 popular datasets and over 700 real-world natural language 18 queries. Our experiments show that Graphy significantly outperforms state-of-the-art natural language based 19 visualization tools, including transformer and rule-based ones. 20

1 INTRODUCTION

1 2

3 4

5

21

22

Natural language interfaces (NLIs) [Gao et al. 2015; Sun et al. 2010; Yu and Silva 2020a] for visualization promise to democratize the visualization authoring process. Given a dataset (often a relational table) and a natural language description, an NLI can generate a set of visualizations that most likely meet the user's intent. For instance, given the dataset shown in Figure 1 and a query such as "give me a scatter plot that shows the fuel economy of all car models", an NLI can, in principle, generate the scatter plot shown on the right side of Figure 1. In this way, even a user with no programming experience can generate visualizations from large-scale data.

While existing NLIs are effective in producing relatively simple visualizations, a recent study [Srinivasan et al. 2021] found that these tools are unable to generate more complex visualizations, such as those that involve subplots or that require performing non-trivial transformations on the input data. For example, for the input query "generate a graph to show the fuel efficiency for cars from different countries segregated based on body style" and the dataset from Figure 1, state-of-the-art tools return the plot shown in Figure 2b as opposed to the ideal plot shown in Figure 2a.

In this paper, we propose a new technique for generating visualizations from natural language descriptions. Our method combines NLP techniques with program synthesis to address several challenging aspects of data visualization. In particular, our technique can handle fairly ambiguous natural language queries, including those that do not fully specify the desired plot type. In addition, our method can perform transformations and aggregations on the input data, allowing it to handle

Authors' addresses: Qiaochu Chen, University of Texas at Austin, Austin, Texas, USA, qchen@cs.utexas.edu; Shankara
 Pailoor, University of Texas at Austin, Austin, Texas, USA, spailoor@cs.utexas.edu; Celeste Barnaby, University of Texas at
 Austin, Austin, Texas, USA, celestebarnaby@utexas.edu; Abby Criswell, University of Texas at Austin, Austin, Texas, USA,
 abbycriswell@utexas.edu; Chenglong Wang, Microsoft Research, Redmond, Washington, USA, chenglong.wang@microsoft.
 com; Greg Durrett, University of Texas at Austin, Austin, Texas, USA, gdurrett@cs.utexas.edu; Isil Dillig, University of
 Texas at Austin, Austin, Texas, USA, isil@cs.utexas.edu.

48 https://doi.org/

^{47 2018. 2475-1421/2018/1-}ART1 \$15.00



Fig. 1. On the left is a cars dataset showing the fuel economy, body style and origin for each model. The plot on the right is for the query "give me a scatter plot that shows the fuel economy of all car models".



Fig. 2. Figures for "show the fuel efficiency for cars from different countries segregated based on body style".

visualizations that require non-trivial data wrangling. As an example, our method can produce the correct plot, shown in Figure 2a, for the input query mentioned earlier.

At the heart of our technique lies a refinement type system that can be used to express properties of the desired visualization. At a high level, our method first uses state-of-the-art NLP techniques, namely a BERT-based [Devlin et al. 2019] intent-and-slots model [Tur et al. 2010], to parse the natural language description into a set of likely refinement type specifications for the desired visualization. Then, for each refinement type specification, our method performs type-directed program synthesis to generate a set of visualization programs of the appropriate type, using a notion of *type compatability* to prune large parts of the search space. Hence, the refinement type system is useful not only as a specification mechanism but also for guiding synthesis and reducing the search space.

A distinguishing feature of the synthesis problem in our setting is that a single visualization task typically results in many synthesis problems, one for each refinement type specification inferred by the parser. However, because each invocation of the synthesizer can be quite expensive, it is important to reuse information across different synthesis problems. Our approach addresses this concern by learning so-called *synthesis lemmas* that can be used to prove unrealizability of future synthesis tasks. In particular, our approach leverages a novel notion of *refinement type interpolants* to learn useful facts that can be reused across different synthesis goals involving the same data set.

We have implemented our proposed approach as a new tool called Graphy and evaluated it on NLVCORPUS [Srinivasan et al. 2021], which consists of 3 popular visualization datasets and over 700 real-world natural language queries. Our evaluation demonstrates that Graphy yields significantly

better results compared to existing state-of-the-art baselines, including transformer and rule-based
 NLIs. We also perform ablation studies to evaluate the importance of our proposed techniques and
 show that they all contribute to making our approach practical

- show that they all contribute to making our approach practical.
 To summarize, this paper makes the following key contributions:
- We propose a new synthesis-based technique for generating visualizations from an input data set and natural language description.
 - We introduce a refinement type system that is useful both as a specification mechanism and for guiding program synthesis.
 - We describe a technique based on the intents-and-slots paradigm for parsing natural language descriptions into refinement type specifications.
 - We propose a type-directed synthesis algorithm that uses a notion of *type compatibility* to prune the search space and learns *synthesis lemmas* that are useful across different synthesis attempts.
 - We implement our approach in a tool called Graphy and perform a large-scale evaluation on over 700 real-world visualization tasks.

2 OVERVIEW

105

106

107

108

109

110

111

112

113 114

115

116

117

118

119 120

121

122

123

124 125

126

We give a high-level overview of our technique with the aid of the motivating example introduced in Section 1. In particular, consider the dataset from Figure 1 listing the fuel economy of different cars and the following natural language query:

"Show the fuel efficiency of cars from different countries segregated based on body style"

Given this query, our tool, Graphy, generates the visualizations shown in Figure 3. Among the plots shown here, the top one is the intended one, with corresponding visualization script shown in Figure 4. We now explain how Graphy is able to generate these visualizations, highlighting salient features of our approach.

Structure of visualization programs. Similar to prior work [Wang et al. 2019], the visualization programs synthesized by Graphy consist of two parts, namely a *table transformation program* P_t





Fig. 4. The visualization program synthesized by Graphy that generates the plot on the top of Figure 3. The top part is a table transformation program that performs a mean operation on the Fuel_economy column. The bottom portion generates a bar chart from the output of the above table transformation program.



and a *plotting program* P_p (see Figure 4). Given these programs, Graphy produces a visualization by first applying P_t to the input data set to obtain a transformed table T and then applying the plotting program P_p to T. Since many real-world visualizations tasks require non-trivial data wrangling, synthesis of table transformations is a crucial aspect of the Graphy workflow. We describe the domain-specific language used for visualizations in more detail in Section 3.

153

171

172

173

174

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

196

Motivation for refinement types. As mentioned in Section 1, Graphy parses the natural 154 language query into a formal specification rather than going directly from natural language to 155 a visualization program. This design choice hinges on two key observations: First, the natural 156 language description often does not contain sufficient information to map it directly to a program. 157 For instance, in our running example, the NL query does not mention anything about a bar 158 graph. Second, the data set to be visualized also contains valuable information for deciding which 159 visualizations make more sense. As an example, looking at the data set, we see that fuel economy 160 is continuous (as opposed to discrete), so it would not be suitable as the x-axis for a bar graph. 161 For these reasons, Graphy parses the NL query into an intermediate specification, which is then 162 supplied as an input to the synthesizer. 163

In this work, we use refinement types as our specification because both base types and logical qualifiers are useful for guiding synthesis. In particular, record types are useful for distinguishing between different types of tabular data, and logical qualifiers capture other forms of hints present in the natural language. For instance, based on the natural language query given above, it is reasonable to conjecture that the color encoding of the plot should be based on country, and our type system allows expressing such information as part of the logical qualifier. We discuss our refinement type system in more detail in Section 4.

From NL queries to refinement types. As a precursor to synthesis, Graphy first uses state-ofthe-art NLP techniques to extract refinement type specifications from the natural language query. In particular, the extracted specifications are of the form $(\mathbb{T}_p, \mathbb{T}_t)$, where \mathbb{T}_p is the output type for the plotting program and \mathbb{T}_t is the output type for the table transformation program. Intuitively, we parse the NL query into two different refinement types as our synthesis procedure generates the plotting and table transformation programs independently.

Our technique for parsing a natural language query to a refinement type consists of two steps. First, we use the technique of *intent classification* [Tur et al. 2010] to infer some of the base types (e.g. BarPlot, ScatterPlot) as well as which *types* of predicates should be involved in the logical qualifier. For our running example, the intent classifier is able to predict that the desired plot is a BarPlot based on the training data. In addition, note that the NL query hints at the fact that the color encoding of the plot is based on country (i.e., "Origin" column in the data set). Such information is encoded using so-called *syntactic constraints* in the logical qualifiers. The intent classifier can decide whether the NL query contains such syntactic constraints, but it cannot decide what the *arguments* of the predicate are. Hence, in a second step, we use a natural-language-processing technique known as *slot filling* [Jeong and Lee 2006] to decide the arguments of the inferred predicates. For our running example, the slot-filling technique can infer that the graph's color is likely to be the Origin field of the input data set and generates a logical qualifier that involves this syntactic constraint. We describe our technique for parsing the NL query into a refinement type specification in more detail in Section 5.

Synthesis workflow. Figure 5 shows the high-level workflow of our synthesis algorithm. For each specification $(\mathbb{T}_p, \mathbb{T}_t)$ generated by the parser, our synthesizer generates a set of visualization programs that satisfy it. At a high level, the synthesis algorithm first generates a plotting program P_p such that the output type of P_p is a subtype of \mathbb{T}_p . The input type of P_p is then used to strengthen

207

208

209

210

213

214

215

217

218

219

220 221

222

223 224

225

226

227

228

229

230

231

232

233

234

235

236

237

238 239

245

197

198

199

200



n_{5 Model} Goal type of n_2 : $\{v : Table(\{Model : \top, Fuel_economy : Qualitative\}) | True\}$ {Model, Fuel_economy}

Fig. 6. Pruning Example. An abstract syntax tree of a partial program is shown on the left. On the right we 211 show the goal type annotation at node n_1 and n_2 . 212

the parsed specification \mathbb{T}_t of the table transformation program to \mathbb{T}'_t . For instance, in our running example, suppose we synthesize the following plotting program:

 $Bar(T, c_x = Body_style, c_y = Fuel_economy, c_{color} = Origin)$

216 Such a program only makes sense if there is a unique y value for every (x, color) pair, so our method strengthens the output of the table transformation program with the following constraint:

 $|Proj(v, \{Body_style, Origin\})| \ge |Proj(v, \{Fuel_economy\})|$

This constraint states that the cardinality (number of unique tuples) of the output table projected on the x (Body_style) and color (Origin) columns should be at least as big as the cardinality when projected onto the y (Fuel economy) column. This constraint serves as a logical qualifier for the output type of the table transformation program and is used to reduce the search space that the synthesizer needs to explore, as discussed in Section 6.1.

Type-directed synthesis. In addition to using refinement types as the specification, our technique also uses them to guide synthesis as in prior work [Polikarpova et al. 2016]. In more detail, our algorithm performs top-down enumerative search, starting with a completely unconstrained program and expanding a non-terminal (i.e., "hole") in the partial program using one of the grammar productions. Each hole is annotated with a so-called goal-type that is propagated backwards using the type system and the initial specification obtained from the NL query. As explained in more detail in Section 6.1, the goal type is used to decide (1) which grammar productions are applicable when performing top-down enumeration, and (2) whether a partial program is infeasible, meaning that its actual type is inconsistent with the annotated goal type. However, unlike prior work, our approach uses a notion of type compatibility (Section 4.3) as opposed to subtyping in order to ensure that we do not rule out correct programs.

For example, suppose we want to synthesize a table transformation that satisfies the following goal type

{
$$v$$
 : Table(Model : Discrete, Fuel_economy : Qualitative) | $\pi(v.Model, count)$ }

for the table in Figure 1. Here, the goal type comes in the form of a refinement type that describes 240 the base type Table(Model : Discrete, Fuel_economy : Qualitative) annotated with the predicate 241 π (v.Model, count). The base type describes a table with attributes Model and Fuel economy, whose 242 types are Discrete and Qualitative respectively. The qualifier, $\pi(v.Model, count)$ is a syntactic 243 constraint that indicates that the Model attribute of the output table is obtained by applying the 244

1:5

1:6 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

count operation. During synthesis, Graphy starts with a program that is a single hole and iteratively 246 expands it using productions in the grammar. Whenever Graphy expands a hole, it propagates 247 the above goal type to newly produced holes in the partial program. For instance, Figure 6 shows 248 a partial program with the annotated goal type of node n_2 . Using our type system, Graphy can 249 prove that this partial program is infeasible because the actual type of the term rooted at node n_2 is 250 incompatible with its annotated goal type. This is because Fuel_economy has type Continuous in 251 the input table, which is inconsistent with the goal type labeling node n_2 , where Fuel_economy is 252 253 required to be Qualitative.

Type-directed lemma learning. A unique feature of our approach is its ability to learn *synthesis lemmas* that can be used across different synthesis tasks involving the same data set.¹ To see why such lemmas are useful, recall that we want to generate multiple visualizations to show to the user, so we need to explore many different programs that *could* be consistent with the NL query. In general, there are multiple plausible specifications one can extract from the NL query, and there are multiple programs that satisfy each specification. Hence, our approach needs to explore many different programs during a single visualization session.

Our approach addresses this concern by using the type system to learn synthesis lemmas. In particular, a *synthesis lemma* is a pair of refinement types (G, R) such that any program with goal type G also needs to be "consistent" (in a sense made precise in Section 4) with refinement type R. Hence, if we encounter a synthesis goal (or sub-goal) that is a subtype of G but that is inconsistent with R, we immediately conclude that the synthesis task is infeasible. Our synthesis algorithm learns such lemmas by inferring so-called *refinement type interpolants* whenever it encounters an infeasible partial program. We discuss the algorithm for type-directed lemma learning in Section 6.2.

Going back to our running example, consider the same infeasible partial program in Figure 6. The root cause of this failure is that the program was unable to convert the Fuel_economy column from a Continuous type to Qualitative. In fact, no program in our DSL can achieve this transformation. Graphy automatically captures this fact by generating the following synthesis lemma:

 $(\{\nu : Table(Fuel_economy : Qualitative)\}, \bot).$

Hence, if we ever encounter a specification such as $\{v : Table(Body_style : Nominal, Fuel_economy : Nominal) | \phi_2\}$ that is a subtype of $\{v : Table(Fuel_economy : Qualitative)\}$, Graphy can immediately conclude that this goal is unrealizable without even attempting synthesis.

278 3 DOMAIN-SPECIFIC LANGUAGE FOR VISUALIZATIONS

In this section, we introduce the domain-specific language for visualization programs. As in prior work [Wang et al. 2019], a visualization program in our setting first performs the necessary table transformations to obtain an intermediate table T and then generates a plot based on T. Hence, as shown in Figure 7, a visualization program P_v can be expressed as the composition of two programs P_t and P_p , where P_t , P_p are programs expressed in the so-called *table transformation* and *plotting* DSLs, respectively. In the remainder of this section, we discuss the syntax and (informal) semantics of these two DSLs in more detail.

Plotting DSL. A program in our plotting DSL takes in an input table T and outputs a plot, which can be one of four types: (1) bar graph, (2) scatter plot, (3) line plot, or (4) area plot. Figure 8 shows an example of each type of visualization supported by our plotting DSL. In more detail, a plotting program is of the form $f(T, c_x, c_y, c_{color}, c_{subplot})$ where f specifies the plot type, T is the input table, and the remaining arguments are attributes of T. Specifically, the c_x , c_y columns specify the

- ²⁹² ¹While there are prior techniques that can learn useful facts during enumeration [Chen et al. 2020; Feng et al. 2018], the facts they learn are not reusable across different synthesis goals.
- 294

254

255

256

257

258

259

260

261

262

263

264

265

266

267

268

269

270

271

272

273

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

Visualization DSL	Sub-DSL for table transformations
$P_v := \lambda T_{in}$. let $T = P_t(T_{in})$ in P_p	$P_t := \lambda T. e$
	e := T
Sub-DSL for plotting $P_p := f(T, c_x, c_y, c_{color}, c_{subplot})$ f := Bar Scatter Line Area	$ bin(e, n, c_{tgt}) $ $ filter(e, val_1 op val_2) $ $ summarize(e, c_{key}, \alpha, c_{tgt}) $ $ mutate(e, c_{tgt}, op, c_{args}) $ $ select(e, c_{args}) $ $val := const c$ $\alpha := mean sum count$

Fig. 7. c denotes column names; const are values in the Table; n is an integer; op is user-provided.



Fig. 8. Examples of plotting programs and their corresponding visualizations.

x- and y-axis of the plot and are required for every program in the plotting DSL. The remaining two arguments c_{color} and $c_{subplot}$ are optional and only make sense for plots with multiple layers or subplots (or both). In particular, the c_{color} attribute is useful for plots that require multiple layers and specifies that each different color in the plot corresponds to a different value of the c_{color} column. Finally, the optional fourth argument specifies that each distinct entry in the $c_{subplot}$ column should be used to generate a different subplot.

Table transformation DSL. As shown in Figure 7, a table transformation program takes in an input table T_{in} , and outputs a table T by applying a sequence of transformations that are inspired by relational algebra and supported by many popular visualization languages, such as VegaLite and ggplot2. In particular, our table transformation DSL includes the following useful constructs:

- The bin operation discretizes a numeric column in the table c_{tgt} into a set of bins. Here, the argument *n* specifies the number of bins that the entries in c_{tgt} should be split into. For example, in the first input table shown in Figure 9, column c_2 is binned.
- The filter construct corresponds to the standard selection operation in relational algebra. Given a table T_{in} and predicate ϕ of the form *val op val*, it produces a subset of T_{in} consisting of all tuples that satisfy ϕ . The second illustration in Figure 9 offers an example of the filter operation.
- The summarize construct performs an aggregation operation specified by α on a specified column c_{tgt} . In more detail, given an input table t and "keys" (i.e., columns) $c_{key} = [c_1, \ldots, c_k]$, it produces a new table that has columns $c_1, \ldots, c_k, c_{tgt}$ such that for each value of the tuple (c_1, \ldots, c_k) , the corresponding value of c_{tgt} is obtained by applying the aggregation operator α to those entries that have the same value for (c_1, \ldots, c_k) . In the third illustration in Figure 9, column c_2 is summarized by the count operator.

344	bin(T, 2, c2)	filter(T, c2 = y1)	summarize(<i>T</i> , [c1], count, c2)	mutate(T, c3, max,[c1, c2])
345 346	c1 c2	c1 c2	c1 c2	c1 c2
347	$\begin{array}{r} x1 + 21 \\ \hline x2 + 27 \end{array}$	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	$\begin{array}{ c c c c c }\hline x1 & y1 \\\hline x2 & y2 \\\hline \end{array}$	$\begin{array}{r} x1 + y1 \\ \hline x2 + y2 \end{array}$
348 349	x1 24	x3 y1	x1 y2	
350		↓ ↓	↓	
351	$\frac{c1}{x1 + 21-25}$	$\begin{array}{c} \mathbf{c1} + \mathbf{c2} \\ \hline \mathbf{x1} + \mathbf{y1} \end{array}$	$\begin{array}{c c} \mathbf{c1} & \mathbf{c2} \\ \hline \mathbf{x1} & 2 \end{array}$	x1 + y1 + max(x1, y1)
352	x2 26-30	x3 y1	x2 1	x2 y2 max(x2, y2)

Fig. 9. Examples of table transformation programs.

- The mutate construct produces a table that has one more column c_{tgt} than its input table. In particular, the value stored in c_{tgt} is obtained by applying operator op to the corresponding values stored in columns $\overline{c_{args}}$. In the fourth illustration in Figure 9, the mutate operator creates column c_3 by taking the max of columns c_1 and c_2 .
- The select construct corresponds to the standard projection operation in relational algebra. In particular, select(t, $\overline{c_{args}}$) yields a table containing only the columns $\overline{c_{args}}$.

Observe that the first argument of each operator is a term *e* in the table transformation DSL; thus, these transformations can be arbitrarily nested within one another. Hence, the table transformation DSL allows performing non-trivial data wrangling tasks that require applying many different operations to the input table.

4 OVERVIEW OF REFINEMENT TYPE SYSTEM

367 While the visualization DSL introduced in Section 3 does not have explicit type annotations, our 368 approach leverages a refinement type system that facilitates effective synthesis. The design of our 369 type system is based on two pragmatic considerations: first, we want our refinement types to serve 370 as useful specifications, meaning that they should capture the clues that are commonly found in 371 natural language descriptions of visualization tasks. Second, we want our type system to be useful 372 for pruning infeasible parts of the search space during synthesis. With these considerations in mind, 373 we introduce those aspects of our refinement type system that are necessary for understanding the 374 overall synthesis approach. 375

4.1 Type Syntax

354

366

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391 392 As standard [Rondon et al. 2008], a refinement type is of the form $\{v : \tau \mid \phi\}$ where τ is a base type and ϕ is a logical qualifier. As shown in Figure 10, base types include strings, integers, four different types of plots, and tables. A table type Table(σ) denotes a table with schema σ , which maps each column name (attribute) to its column type, which indicates the type of values stored under that column. The column type \top indicates *any* type of of data, whereas Quantitative and Qualitative indicate whether the entry is associated with a quantity or quality respectively. Quantitative data can be further divided into Continuous and Discrete, and Qualitative data can be divided into Nominal, Ordinal, and Temporal (e.g., date or year).

In contrast to base types, logical qualifiers are formulas formed from atomic predicates using the standard logical connectives \neg , \land , and \lor . We differentiate between two types of atomic predicates, namely *syntactic constraints* $\pi(..)$ and *table property predicates* of the form $\theta \notin \theta$. We discuss both types of atomic predicates in more detail below.

Syntactic constraints. Given a table or plot *x* with attribute *c*, the predicate $\pi(x.c, \mu)$ expresses that μ was used in the derivation of *x.c.* Here, μ is either a built-in function *f* in our DSL (e.g.,

393	Base Type	Refinement Type
394	$\tau := \text{Table}(\sigma) \mid \tau_p \mid \text{str} \mid \text{int}$	$\mathbb{T} := \{ \nu : \tau \mid \phi \} \mid x : \mathbb{T} \to \mathbb{T}$
395	$\tau_p := BarPlot ScatterPlot$	$\phi := \pi(x.\eta, \mu) \mid \theta \ \emptyset \ \theta \text{ where } \emptyset \in \{=, \ge, \le\}$
396	LinePlot AreaPlot	$ \neg \phi \phi \land \phi \phi \lor \phi$
397	$\sigma := \{c_1 : i_{c_1},, c_n : i_{c_n}\}$ $\tau_c := \top \text{Oualitative} \text{Ouantitative} $	$\eta := x y \text{color} \text{subplot} c$ u := mean sum count bin filter mutate x c
398	Nominal Ordinal Temporal	$\theta := \gamma \rho(\gamma) n c x$
399 400	Discrete Continuous	$\gamma := x \operatorname{Proj}(\gamma, \overline{c}) \operatorname{Filter}(\gamma, val_1 op val_2)$
400	Fig. 10. Type Syntax. <i>c</i> is a column name; <i>const</i> ar	$\rho := \max \min$ e values in the Table; <i>n</i> is an integer; <i>x</i> is a variable.
402 403 404 405 406	count, mutate) indicating that function f was in form $x'.c'$ indicating data flow from $x'.c'$ to $x.c.$ system allow encoding useful hints present in the data used in the visualization task.	volved in the computation of $x.c$, or a term of the Intuitively, the syntactic constraints in our type e natural language description about the origin of
407 408	Example 4.1. For the running example from Sect the output of the plotting program:	ion 2, our NL parser generates the following type for
409 410	$\{v: BarPlot \mid \pi(v)\}$	v.color, x.Origin)}
411	where x refers to the input of the plotting program	n. Here, the base type indicates that we want a bar
412	graph, and the syntactic constraint indicates that it	the color encoding is bound to the Origin field of the
413	input table. In other words, it indicates that colors i	n the plot correspond to values of the Origin column.
414	Table properties. In addition to the syntactic	requirements our type system allows expressing
415	properties of tables using predicates of the form	$\theta \land \theta$ where θ is a term and δ is a relation symbol
417	(e.g., \leq). In more detail, terms θ can be formed u	sing the following constructs:
418	• Given a variable x of type Table, $ x $ represents	the cardinality of <i>x</i> (i.e., number of unique tuples).
419 420	• The functions Proj and Filter have the same s	semantics as the corresponding constructs in our
421 422	 Given a column <i>x</i>, the aggregation operators minimum) value in <i>x</i>. 	s $max(x)$ and $min(x)$ return the maximum (resp.
423 424	Intuitively, table property predicates are useful of the visualization task and provide significant	for specifying the table transformation component pruning power during synthesis.
425	Example 4.2. Consider the following refinement	type:
427	$\{v: Table(Price : Discrete, Origin : Nom\}$	inal) $ v = 3 \land \max(\operatorname{Proj}(v, \{\operatorname{Price}\})) = 8\}$
428	This type describes a table that (1) has two attribu	tes, Price and Origin, of types Discrete and Nominal
429	respectively, (2) contains three unique tuples, and	(3) has a maximum value of 8 in its Price column.
430 431	4.2 Subtyping	
432	Given a refinement type specification \mathbb{T} the gas	al of our approach is to synthesize a visualization
433	program of type \mathbb{T}' such that \mathbb{T}' is a subtype of	f T. Thus, we start by formalizing the subtyping
434	relation for our type system using judgments of	the following form:
435	$\Gamma \vdash \mathbb{T}$	$_{1} <: \mathbb{T}_{2}$
430 437	where Γ is a type environment manning variable	s (and built-in DSL functions) to their correspond-
438	ing types. As standard, the meaning of this iuc	Igment is that \mathbb{T}_1 is a subtype of \mathbb{T}_2 under type
439	environment Γ . Since deciding subtyping between	a base types does not require the type environment,
440	we omit the type environment for base types.	
441		

1:9

 \vdash Qualitative <: \top

Base-Trans $\frac{\vdash \tau'' <: \tau' \vdash \tau' <: \tau}{\vdash \tau'' <: \tau} \qquad \text{Base-Ref} \vdash \{v : \tau \mid \phi\} <: \tau$

TABLE-WIDTH \vdash Table({ $c_i : \tau_i \stackrel{i \in 1...n+k}{\rightarrow}$ }) <: Table({ $c_i : \tau_i \stackrel{i \in 1...n}{\rightarrow}$ })

 $\text{Table-Depth} \quad \frac{\forall i. \vdash \tau_i <: \tau'_i}{\vdash \text{Table}(\{c_i : \tau_i^{i \in 1...n}\}) <: \text{Table}(\{c_i : \tau_i^{\prime i \in 1...n}\})}$

TABLE-PERMUTATION $\frac{\vdash \mathsf{Table}(\{c_i : \tau_i \ ^{i \in 1...n}\}) \text{ is a permutation of } \mathsf{Table}(\{c'_i : \tau_i \ ^{i \in 1...n}\})}{\vdash \mathsf{Table}(\{c_i : \tau_i \ ^{i \in 1...n}\}) <: \mathsf{Table}(\{c'_i : \tau_i \ ^{i \in 1...n}\})}$

⊢ Discrete <: Quantitative

+ Temporal <: Qualitative

⊢ Ordinal <: Qualitative

 $\operatorname{Ref} \frac{\vdash \tau_1 <: \tau_2 \qquad \operatorname{Encode}(\Gamma) \land \operatorname{Encode}(\phi_1) \Rightarrow \operatorname{Encode}(\phi_2)}{\Gamma \vdash \{\nu : \tau_1 \mid \phi_1\} <: \{\nu : \tau_2 \mid \phi_2\}} \qquad \operatorname{Func} \frac{\Gamma \vdash \mathbb{T}'_1 <: \mathbb{T}_1 \quad \Gamma \vdash \mathbb{T}_2 <: \mathbb{T}'_2}{\Gamma \vdash x : \mathbb{T}_1 \rightarrow \mathbb{T}_2 <: x : \mathbb{T}'_1 \rightarrow \mathbb{T}'_2}$

Ś

 \vdash Quantitative <: \top

⊢ Continuous <: Quantitative</p>

⊢ Nominal <: Qualitative</p>

Fig. 11. Base and refinement type subtyping relation.

Symmetry
$$\frac{\vdash \tau \sim \tau'}{\vdash \tau' \sim \tau}$$
 Data $\frac{\tau, \tau' \in \tau_c \quad \vdash \tau <: \tau' \lor \vdash \tau' <: \tau}{\vdash \tau \sim \tau'}$

$$\text{TABLE} \frac{\forall i, j. (c_i = c'_j) \rightarrow \vdash (\tau_i \sim \tau'_j)}{\vdash \text{Table}(\{c_i : \tau_i^{i \in 1...n}\}) \sim \text{Table}(\{c'_j : \tau'_j^{j \in 1...m}\})} \qquad \text{Func} \frac{\Gamma \vdash \mathbb{T}_1 \sim \mathbb{T}'_1 \quad \vdash \mathbb{T}_2 \sim \mathbb{T}'_2}{\Gamma \vdash x : \mathbb{T}_1 \rightarrow \mathbb{T}_2 \sim x : \mathbb{T}'_1 \rightarrow \mathbb{T}'_2}$$

 $\begin{array}{l} \vdash \tau_{1} \sim \tau_{2} \\ \\ \text{Refinement-Comp} \end{array} \frac{\mathsf{SAT}(\mathsf{Encode}(\Gamma) \land \mathsf{Encode}(\phi_{1}) \land \mathsf{Encode}(\phi_{2}))}{\Gamma \vdash \{\nu : \tau_{1} \mid \phi_{1}\} \sim \{\nu : \tau_{2} \mid \phi_{2}\}} \end{array}$

Fig. 12. Base and refinement type compatibility relation

Figure 11 presents our subtyping rules. The first several rules are straightforward and show the subtyping relation between primitive types like Discrete and Quantitative. The subtyping rules for tables are essentially standard subtyping rules for records [Pierce 2002]. The last two rules for refinement types are also standard and require (1) checking the subtyping relation between base types and (2) checking the validity of a logical formula for the logical qualifiers. In particular, these rules make use of a function called Encode that converts the logical qualifier of a refinement type into an SMT formula. The interested reader can find details of the SMT encoding in the appendix.

4.3 Type compatibility

While our synthesis algorithm ensures that the type of the synthesized program is a subtype of the specification, we utilize a weaker notion of *type compatibility* for pruning during synthesis. In particular, because our synthesis algorithm needs to reason about the feasibility of incomplete programs (where some parts are yet to be determined), we introduce a notion of type compatibility that is much weaker than subtyping. Intuitively, two types \mathbb{T}_1 and \mathbb{T}_2 are *compatible* with each other if there exists a subtype \mathbb{T} of \mathbb{T}_1 that is also a subtype of \mathbb{T}_2 . Conversely, if two types \mathbb{T}_1 and \mathbb{T}_2 are incompatible, there is no refinement of \mathbb{T}_1 that will make it a subtype of \mathbb{T}_2 . As we will see in Section 6, the notion of type (in)compatibility is very useful for pruning during synthesis. In this section, we formalize this notion and present rules for checking type compatibility. We define the

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

515 516

517

521

Type-Directed Synthesis of Visualizations from Natural Language Queries

$$\Gamma(T) = \{v : \tau_T \mid \phi_T\}$$

$$\tau_T = \text{Table}(\{c_x : \text{Discrete}, c_y : \text{Quantitative}, c_{\text{color}} : \text{Discrete}, c_{\text{subplot}} : \text{Discrete}\})$$

$$BAR \frac{\text{Encode}(\Gamma) \land \text{Encode}(\phi_T) \Rightarrow |(v, \{c_x, c_{\text{color}}, c_{\text{subplot}}\})| \ge |(v, \{c_y\})|}{\Gamma \vdash \text{Bar}(T, c_x, c_y, c_{\text{color}}, c_{\text{subplot}}) : \{v : \text{BarPlot} \mid \bigwedge_{e \in \{x, y, \text{color}, \text{subplot}\}} \pi(v.e, T.c_e)\}}$$

Fig. 13. Typing Rule for a Bar Plot. We use notation $(v, \{c_1, \ldots, c_n\})$ as a shorthand for $Proj(v, \{c_1, \ldots, c_n\})$.

compatibility relation for our type system using judgments of the form:

 $\Gamma \vdash \mathbb{T}_1 \backsim \mathbb{T}_2$

stating that \mathbb{T}_1 is compatible with \mathbb{T}_2 under environment Γ , as shown in Figure 12. Unlike the subtyping relation, the compatibility relation is symmetric (first rule in Figure 12); however it is not transitive. The second rule in Figure 12 defines the type compatibility relation for primitive types and states that they are compatible if one is a subtype of the other or vice versa. The third rule (TABLE) asserts that two table types are compatible when all their shared columns are compatible. The intuition is that if all shared columns are type compatible, then we can construct a new table type that is a refinement of both by taking the union of their schemas. Finally, REFINEMENT-COMP and FUNC are similar to their subtyping counterparts in that they reduce the compatibility check to an SMT query. However, there are two key differences. First, the encoded formula is a conjunction of the qualifiers as opposed to an implication. Second, we check that the encoding is satisfiable as opposed to valid. The intuition behind this rule is that, if the resulting formula is satisfiable, then $\{v: \mathbb{T} \mid \phi_1 \land \phi_2\}$ is a well defined type in our type system that has at least one inhabitant, and it refines both $\{v : \mathbb{T} \mid \phi_1\}$ and $\{v : \mathbb{T} \mid \phi_2\}$.

4.4 Typing Rules

In this section, we give an overview of our typing rules for assigning types to DSL terms. In 518 particular, our typing rules derive judgments of the form $\Gamma \vdash t : \mathbb{T}$ to indicate that term t has type 519 T under environment Γ . Since the typing rules are not the primary focus of this paper, we only 520 discuss two representative rules and leave the rest to the appendix.

522 Typing rules for the plotting sub-DSL. To illustrate the typing rules for the plotting sub-DSL, 523 Figure 13 shows the rule for the Bar construct, which generates a bar graph given table T. At a 524 high level, this rule states that if T's type satisfies two constraints, then the output type will be a 525 refinement of BarPlot. The first constraint is that T's schema must be suitable for generating bar 526 graphs, meaning that c_x is Discrete and c_y is Quantitative. This requirement is captured by the 527 second premise. In addition to having a suitable schema, another important requirement for a bar 528 graph is that it should not have overlapping bars, meaning that the x-label in each subplot must 529 correspond to a unique y-value. This requirement is captured through the cardinality constraint 530 in the third premise, which checks that the logical qualifier for T implies that there is unique y531 for each x. If these premises hold, then the entire term is well-typed with base type BarPlot and a 532 logical qualifier stating the syntactic constraint for the return value of Bar. In particular, the logical 533 qualifier in the conclusion states, for example, that the x attribute of the plot is derived from the c_x 534 attribute of the input table *T*.

Typing rules for table transformation sub-DSL. Figure 14 shows the typing rule for the summarize construct in our table transformation DSL. Recall that summarize takes as input an aggregation operator, and the type depends on which aggregation operator summarize is invoked with. In

538 539

535

536

1:12 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

540 541 542

$$\Gamma \vdash e : \{v : \tau_t \mid \phi\} \text{ where } \tau_t = \mathsf{Table}(\{\dots, \mathsf{c}_{\mathsf{tgt}} : \tau_{\mathsf{tgt}}, \dots\})$$

$$c_{tgt} \notin c_{key} \vdash \tau_{tgt}$$
: Quantitative

 $\tau' = \text{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k, c_{\text{tgt}} : \tau_{\text{tgt}}\}) c'_i \in \overline{c_{\text{key}}} \quad \tau' = \tau'[c_{\text{tgt}} \mapsto \text{Continuous}]$

$$\phi_1 = \phi \notin \text{Terms}(\phi, c_{\text{tgt}})$$
 $\phi_2 = \phi_1 \notin \pi(v.c_{\text{tgt}}, \text{mean})$

543 544 545

 $\frac{\phi' = \phi_2 \land |(v, \{c_{tgt}\})| \le |(v, \overline{c_{key}})| \land \pi(v.c_{tgt}, \text{mean})}{\Gamma \vdash \text{summarize}(e, \overline{c_{key}}, \text{mean}, c_{tgt}) : \{v : \tau' \mid \phi'\}}$ Summ-Mean —

546 Fig. 14. Typing rule for Summarize instantiated with a Mean operation. We use notation $(v, \{c_1, \ldots, c_n\})$ as a 547 shorthand for $Proj(v, \{c_1, ..., c_n\})$, and the \notin operator is defined in the text. 548

Figure 14, we consider instantiating summarize with mean as a representative example. To under-549 stand this typing rule, let us first recall the semantics of summarize, which associates each unique 550 value of the specified key columns with the mean of the values in the specified target column (see 551 552 Figure 9). The first two premises in the typing rule SUMM-MEAN impose some requirements on the input table. In particular, because it only makes sense to take the mean of quantitative values, 553 the second premise ensures that the target column has a suitable type. Furthermore, since the 554 mean operation produces a value of type Continuous, the column c_{tat} has type Continuous in 555 the output table with base type τ' . The fourth and fifth lines in Figure 14 state the relationship 556 between the logical qualifiers of the input and output tables. To that end, given a logical qualifier 557 ϕ and a set of terms *S*, we use the notation $\phi \notin S$ to denote the strongest logical qualifier ϕ' that 558 is implied by ϕ and that does not imply anything about any term $t \in S$.² Thus, according to our 559 typing rule, the new logical qualifier ϕ' for the output table differs from the qualifier ϕ for the 560 input table in the following ways: First, it "removes" from ϕ any knowledge about the terms that 561 562 involve c_{tgt} which are affected by the summarize operation. Second, it asserts that the number of unique tuples over $\overline{c_{key}}$ is greater than or equal to the number of unique values in c_{tgt} . This 563 is because the cardinality of the output table is equal to the number of unique (c_1, \ldots, c_k) values 564 where each $c_i \in \overline{c_{keys}}$. However, as two distinct (c_1, \ldots, c_k) values could have the same value for 565 $c_{\rm tgt},$ we cannot infer a stronger constraint. Finally, since the values of $c_{\rm tgt}$ were produced by the 566 mean operation, ϕ' includes the syntactic constraint $\pi(v.c_{tgt}, mean)$. 567

FROM NATURAL LANGUAGE TO REFINEMENT TYPES 5

In this section, we describe a technique for generating refinement type specifications from natural 570 571 language queries. At a high level, we frame this problem as an instance of the intents-and-slots 572 problem [Jeong and Lee 2006; Tur et al. 2010] and build a parser that combines intent detection and 573 slot filling on top of the BERT language model [Devlin et al. 2019].

Background on Intents-and-Slots-Paradigm 575 5.1

576 The intents-and-slots paradigm is a classical paradigm in the NLP literature on task-oriented dialog 577 systems [Dahl et al. 1994; Hemphill et al. 1990; Tur et al. 2010] and flexibly supports many types of 578 user interactions, as evidenced by its adoption on Amazon Alexa and other dialog platforms. At a 579 high level, *intent classification* is the problem of determining the topic of a query from a natural 580 language utterance. For instance, given a set of topics such as "flights", "movies", "restaurants", 581 intent classification can be used to determine which of these topics a sentence is about. Once the 582 intent of the utterance is identified, *slot filling* determines pre-defined properties of that topic. For 583 example, if the topic of a query is "flights", relevant parameters include airline, destination city, 584 flight number, etc., and slot-filling techniques aim to identify these parameters.

588

585

568

569

⁵⁸⁶ ²One way to obtain ϕ_{z}^{t} is to replace all occurrences of t with a fresh existentially quantified variable x and then eliminate the quantifier. We formalize the $\frac{1}{2}$ operator in the appendix. 587



Fig. 15. Model architecture for the color encoding property in the running example. The yellow box represents the contextualized embedding of the query. Blue boxes represent the BERT embedding for each of the columns names in the table. After applying a linear layer to the BERT encoding in each task, we obtain a probability distribution across all possibles classes. The one with the highest probability is highlighted with pink.

As a concrete example, consider the query "What flights are available from San Francisco to New York?" Here, an intent classifier aims to determine that the topic of the query belongs to the flight category as opposed to movies or restaurants. Then, assuming that the flight category has attributes such as departure and destination city, a slot-filling technique can be used to determine that the departure city of the query is San Francisco and that the destination is New York.

The intents-and-slots paradigm is a good fit for our setting for two main reasons. First, compared to conventional semantic parsing [Zelle and Mooney 1996; Zettlemoyer and Collins 2005], the intents-and-slots framework does not make strict assumptions about the grammatical structure of inputs. As a result, it can more flexibly handle user inputs that do not conform to a pre-defined syntax. Second, user queries in our setting can be naturally classified into different intents based on (1) the type of the plot (e.g., bar graph, scatter plot, etc.) they refer to and (2) which predicates in our refinement type system they involve. Furthermore, the arguments of these predicates can be determined using the slot-filling paradigm.

5.2 Parsing Technique

In this section, we explain our instantiation of the intents-and-slots paradigm for our setting.

Overview. We have identified six types of properties that are typically mentioned in natural language queries and that are useful to the synthesizer. These include the following:

- Plot type: According to the query, what is the most likely plot type desired by the user?
- Color: Does this query mention anything about the color encoding of the plot?
- Subplot: Does the query indicate that the visualization has subplots?
- Mean: Does the query indicate that the visualization requires computing the mean of values?
- **Sum:** According to the query, does the visualization require summing values?
- Count: Does the query indicate that the visualization involves the use of the count operator?

Our parser uses six different intent classifiers, one for each category listed above. With the exception of plot type, all intent classifiers are binary and yield a yes/no prediction. In the case of a "yes" prediction, our parser uses slot filling to predict which attribute in the source data set this operation is associated with. For plot types, the intent classifier predicts whether the query refers to a bar chart, scatter plot, line graph, or area plot.

Input to BERT. Our method performs the predictions outlined above using a fine-tuned BERT model, with a shared encoder for both intent classification and slot filling as illustrated in Figure 15.

The input to BERT is of the following form:

638 639

640

656

657

658

659

660

661

662

663 664 665

666

667

668

669

670

671

[CLS] N [SEP] c_1, \ldots, c_n

where *N* is the natural language query, c_1, \ldots, c_n denote column names from the input table *D*, and 641 [CLS] and [SEP] are standard placeholder tokens. We include the column names as part of the 642 input for two reasons. First, when performing slot filling, the model needs to predict attributes of 643 the source table, so we will need access to embedded representations of the column names. Second, 644 even for intent classification, information about the input table provides useful context that the 645 BERT model can condition on. Given this input, BERT generates a contextualized encoding of each 646 of its input tokens, where each token in the input is mapped to a dense vector representation 647 informed by all the other tokens through BERT's attention mechanism [Vaswani et al. 2017]. 648

Intent Classification. Our intent classifier is a model of the form $p(C_i | N, A_D; \mathbf{w}_{c_i})$, where A_D is the set of column names for input table D, and $C_i \in \{0, 1\}$ is a binary label for classifiers other than plot type, and $C_i \in \{\text{Bar, Scatter, Line, Area}\}$ for the plot type classifier. Additionally, N denotes the NL input, i is an index denoting the property type, and w_i are the model weights. As standard practice, we take the vector \mathbf{h}_{CLS} as the representation of the sentence, and we use $p(C_i | N, A_D, i; w_i) = \sigma(\mathbf{w}_{c_i}^\top \mathbf{h}_{\text{CLS}}(N, A_D))$, where σ denotes the logistic function, making this a standard logistic regression layer with the CLS token's contextualized embedding as input.

Slot-filling model. As shown in Figure 15, a property like color requires a parameter in the form of one of the column names. Crucially, such arguments cannot be predicted with a standard classification model since the column names change with each table D being plotted. As such, the model must be able to place a distribution over an arbitrary set of column tokens. To address this issue, we use a pointer mechanism similar to implementations of the attention mechanism in settings like machine translation [Bahdanau et al. 2014] and document summarization [See et al. 2017]. Specifically, our slot-filling model places a distribution $p(c_i | N, A_D, i; w_i)$ over column names. We use the same BERT encodings as the intent classifiers and let

$$p(c_i \mid N, A_D, i; \mathbf{W}) = \operatorname{softmax}_i(\mathbf{h}_{\operatorname{CLS}}(N, A_D)^{\top} \mathbf{W} \mathbf{h}_{c_i}(N, A_D))$$

where **W** is a square weight matrix and softmax denotes the standard softmax operation, which exponentiates and normalizes the arguments to form a probability distribution.

From BERT predictions to specifications. Recall that the input to our synthesizer is a pair of refinement types of the form $(\mathbb{T}_p, \mathbb{T}_t)$ where \mathbb{T}_p is the output type of the plotting program and \mathbb{T}_t is the output type of the table transformation program. We now explain how to map the predictions made by the BERT model to specifications of this form.

To generate \mathbb{T}_p for the plotting program, we assign the prediction of the plot type classifier to be the base type of \mathbb{T}_p . The qualifier of \mathbb{T}_p consists of a conjunction of syntactic constraints output by the color and subplot models. In particular, the logical qualifier of \mathbb{T}_p includes a syntactic constraint $\pi(v.color, x.f)$ if the intent classifier for the color property predicts "yes" and the slot-filling model outputs column name f.

In particular, for each model, we generate the syntactic constraint for this specific property if the intent classifier returns "yes" and populate the arguments of the predicate with the output of the slot-filling model. For instance, in Figure 15 where we focus on the color property, we first generate the predicate template $\pi(v.color, x.?)$ where ? is to be determined by the argument classifier. Then, ? is filled by the output of the argument classifier, which in this case is Origin and the model returns the predicate $\pi(v.color, x.Origin)$ as its final output.

The base type of \mathbb{T}_t is obtained using a pre-trained model [Lin et al. 2020] that outputs the set of likely columns in the table mentioned in the query. For the logical qualifier of \mathbb{T}_t , we use the predictions made by the mean, sum and count intent classifiers and their corresponding slot-filling

models. For example, the logical qualifier includes a predicate $\pi(v.f, \text{sum})$ if the sum model predicts "yes" and the slot-filling mechanism predicts column name f for the first argument.

⁶⁸⁹ Distribution over specifications. Our parser assigns a probability to each specification by using the ⁶⁹⁰ probabilities output by each model. In particular, let us view a refinement type \mathbb{T} as a set of tuples ⁶⁹¹ { C_i, c_i }_i where C_i is a intent and c_i is an attribute predicted by the slot-filling model. Then, our ⁶⁹² method assigns a probability to these sets of tuples as $p({C_i, c_i}_i) = \prod_i p(C_i | N, A_D, i; \mathbf{w}_{c_i})p(c_i | N, A_D, i; \mathbf{W})$. Hence, we can rank all possible specifications from highest to lowest probability.

6 SYNTHESIS FROM REFINEMENT TYPE SPECIFICATIONS

In this section, we describe our synthesis algorithm which takes as input a visualization specification 696 $(\mathbb{T}_p, \mathbb{T}_l)$ and an input table *D* and generates *all* visualization programs $P_v = P_p \circ P_l$ such that 697 (1) $P_t(D)$ is an inhabitant of \mathbb{T}_t (written $P_t(D) \models \mathbb{T}_t$) and (2) $P_v(D)$ is an inhabitant of \mathbb{T}_p . At a 698 high level, the synthesis algorithm is based on type-directed top-down enumerative search and 699 uses the refinement type system from Section 4 to significantly reduce the search space. We first 700 start by explaining the basic synthesis algorithm (Section 6.1) and then introduce the concept of 701 type-directed lemma learning to improve the scalability of our approach (Section 6.2). Our algorithms 702 frequently use the typing judgements from Section 4. While these typing judgments make use of a 703 type environment, we treat the type environment as implicit and drop it to simplify presentation. 704

706 6.1 Overview of Synthesis Algorithm

694

695

705

721

722

723

724

731

732

733

734 735

Our top-level synthesis algorithm is presented in Figure 16a and works as follows: Given a table 707 D of type \mathbb{T}_{in} and specification $(\mathbb{T}_p, \mathbb{T}_t)$, it first synthesizes a set of plotting programs \mathcal{P}_p whose 708 output type is a subtype of the goal type (line 3). In more detail, each plotting program $P_p \in \mathcal{P}_p$ of 709 type $\mathbb{T}_{in}^{p} \to \mathbb{T}_{out}^{p}$ satisfies the following two properties: (1) $\mathbb{T}_{out}^{p} <: \mathbb{T}_{p}$ and (2) $\mathbb{T}_{in}^{p} \sim \mathbb{T}_{t}$. The first 710 constraint ensures that the generated visualization satisfies the user's specification, and the second 711 constraint ensures that there is at least one input table to the plotting program that is consistent with 712 \mathbb{T}_t . Then, for each synthesized plotting program P_p of type $\mathbb{T}_{in}^p \to \mathbb{T}_{out}^p$, the algorithm synthesizes (at line 6) a set of corresponding table transformation programs \mathcal{P}_t of type $\mathbb{T}_{in}^t \to \mathbb{T}_{out}^t$ such 713 714 that (1) $\mathbb{T}_{out}^t \sim \mathbb{T}_t \wedge \mathbb{T}_{in}^p$ and (2) $\mathbb{T}_{in} <: \mathbb{T}_{in}^t$. Note that the first condition *strengthens* the original 715 specification using \mathbb{T}_{in}^{p} (via intersection types) and ensures that there is at least one output of the 716 717 table transformation program that is a valid input to the plotting program.

The key part of the algorithm is the SYNTHESIZEGOAL procedure, presented in Figure 16c, that
 is used to synthesize *both* table transformation and plotting programs. To unify presentation,
 SYNTHESIZEGOAL takes a few additional arguments:

- \mathcal{G} , the grammar for the DSL in which we synthesize programs
- The correctness checking condition \triangleright for the *input* type (either <: or \backsim)
- The correctness checking condition ⊲ for the *output* type (either <: or ∽)

At a high level, SYNTHESIZEGOAL is a top-down enumeration procedure which starts from the root symbol of the grammar and keeps expanding non-terminals until it generates a complete program. We represent the syntax of the underlying DSL as a context-free grammar $\mathcal{G} = (V, \Sigma, R, S)$, where V, Σ denote non-terminals and terminals respectively, R is a set of productions, and S is the start symbol. As standard [Feng et al. 2018], we formalize our top-down enumeration procedure using the notion of *partial programs*:

Definition 6.1 (Partial program). A partial program P is a sequence $P \in (\Sigma \cup V)$ * such that $S \stackrel{*}{\Rightarrow} P$ (i.e. P can be derived from S via a sequence of productions). We refer to any non-terminal in P as a hole, and we say that P is complete if it does not contain any holes. 1:16 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

In the remainder of this section, we represent each partial program P as an abstract syntax tree 736 (AST) (*N*, *E*) with nodes *N* and edges *E*. Each node $n \in N$ is represented as a pair (l, \mathbb{T}_{\perp}) where 737 *l* is a node label (either a terminal or non-terminal symbol in *G*) and \mathbb{T}_{\perp} is the *goal type* of the 738 subprogram rooted at *n*. The goal type \mathbb{T}_{\perp} of a node *n* serves as a necessary correctness condition 739 such that if the sub-program rooted at *n* does not satisfy \mathbb{T}_{\downarrow} , then the whole program cannot satisfy 740 its specification. For a node n, we use the notation P(n) to denote the subtree of P rooted at n, and 741 use Label(n) and GoalType(n) to refer to the label and goal type of n, respectively. Finally, we refer 742 743 to a node as *complete* if the subtree rooted at *n* is a complete program.

With this notation in place, we now describe the basic version (everything not underlined) of 744 SYNTHESIZEGOAL in more detail. Our algorithm maintains a worklist W of partial programs and 745 iteratively grows it. At the beginning, W is initialized to be the empty program P_0 with a single 746 node n_0 annotated with the grammar start symbol S_G and top-level goal \mathbb{T}_{out} . The loop in lines 747 4-18 dequeues a program P from the worklist with type $\mathbb{T}_{in}^P \to \mathbb{T}_{out}^P$ and checks if it is complete and 748 whether it satisfies the correctness conditions. If so, this program is added to the set S containing 749 all synthesis results. Otherwise, SYNTHESIZEGOAL calls Expand at line 12 to generate a new set of 750 partial programs by expanding a hole h in P. Similar to prior work [Feser et al. 2015; Polikarpova 751 et al. 2016], when Expand generates a new partial program P', it propagates the goal type at h to 752 753 its children. However, the goal types we infer are necessary conditions for correctness with respect to type compatibility (as opposed to subtyping) and are derived based on the premises of the typing 754 rules from Section 4.4. In other words, the types of all subprograms must be compatible with their 755 propagated goal type in order for the overall program to be compatible with its goal type. 756

Next, for each expansion P' of P, the TYPEINCOMPATIBLE procedure (presented in Figure 16b) 757 uses our refinement type system to check whether P' is infeasible. To do so, it iterates over all 758 nodes and checks whether the subtree rooted at that node is a complete program (line 3). If so, 759 it infers the type \mathbb{T} of this sub-program using our type system (line 4) and queries whether \mathbb{T} is 760 type-compatible with the goal type of *n*. 761

THEOREM 1. Let P be a partial program with input type \mathbb{T}_{in} and top level goal type \mathbb{T}_{out} . If *TYPEINCOMPATIBLE*(*P*) returns true, then for any completion *P'* of *P*, *P'* \neq ($x : \mathbb{T}_{in} \rightarrow \mathbb{T}_{out}$).

6.2 Type-Directed Learning

We now describe our *type-directed learning* technique that refines the basic synthesis algorithm from the previous subsection. The motivation for this technique is that we need to synthesize many programs during each visualization session. To leverage the similarities across all these synthesis tasks, our algorithm learns so-called synthesis lemmas that capture inferred constraints for the input data set. While this idea is somewhat similar to the notion of conflict-driven learning in prior synthesis work [Feng et al. 2018], there are two key differences. First, our learned lemmas can be reused across different specifications as long as the input data set is the same. Second, the learning of the synthesis lemmas is type-directed and leverages our refinement type system. 775

776 **Definition 6.2 (Synthesis lemma).** A synthesis lemma for an input table D is a pair of refinement 777 types (G, R) such that, for any program P and type \mathbb{T} satisfying $\mathbb{T} <: G$, if P(D) is an inhabitant of 778 **T**, then we have $\mathbb{T} \sim R$. 779

In other words, a synthesis lemma captures additional (learned) constraints R that the synthesized 780 program must satisfy if its output type is to be a subtype of G. Given a lemma (G, R) and partial 781 program *P*, the basic idea is to use *R* for pruning as follows: If the desired goal type \mathbb{T} of *P* is a 782 subtype of G but \mathbb{T} is *not* type-compatible with R, then we can prune P without even attempting 783

784

762

763

764 765 766

767

768

769

770

771

772

773

785			
786	1: procedure SynthesizeVis($(\mathbb{T}_p, \mathbb{T}_t), D$)		
787	input: A specification $(\mathbb{T}_t, \mathbb{T}_p)$	1:	procedure SynthesizeGoal($\mathcal{G}, \mathbb{T}_{in}, \mathbb{T}_{out}, \triangleright, \triangleleft$)
788	input: The input table <i>D</i>		input: Crammar G Specification (\mathbb{T}, \mathbb{T})
789	output: A set of visualization programs.		input: Operators $\triangleright \triangleleft \in \{c: a\}$ to check correct
790	2: $\mathcal{S} \leftarrow \emptyset; \mathbb{T}_{in} \leftarrow \text{GetType}(D)$		ness for input and output type respectively
791	3: $\mathcal{P}_p \leftarrow \text{SynthesizeGoal}(\mathcal{G}_v, \mathbb{T}_t, \mathbb{T}_p, \sim, <:)$	_	
792	4: for all $P_p: \mathbb{T}_{in}^p \to \mathbb{T}_{out}^p \in \mathcal{P}_p$ do	2:	$\mathcal{S} \leftarrow \{\}$
793	5: $\mathbb{T}_{s} \leftarrow \mathbb{T}_{t} \wedge \mathbb{T}_{in}^{p};$	3:	$P_0 \leftarrow \{(S_G, \mathbb{1}_{out}), \emptyset\}; W \leftarrow \{P_0\}$
794	6: $\mathcal{P}_t \leftarrow \text{SynthesizeGoal}(\mathcal{G}_t, \mathbb{T}_{in}, \mathbb{T}_s, <:, \sim)$	4:); <u>5</u> .	while $W \neq \emptyset$ do $P \leftarrow W remove()$:
795	7: for all $P_t \in \mathcal{P}_t$ do	- J. 6.	$T \leftarrow W$.remove(), $\mathbb{T}^P \leftarrow \text{InputType}(P)$
796	8: if $P_t(D) \models \mathbb{T}_t \land P_p(P_t(D)) \models \mathbb{T}_p$ then	7.	$\mathbb{T}_{in}^{P} \leftarrow \text{OutputType}(P)$
707	9: $\mathcal{S} \leftarrow \mathcal{S} \cup \{P_v \circ P_t\};$	7. 8.	$\mathbf{m}_{out} \leftarrow \text{Output type}(\mathbf{r})$ if is Complete(P) then
709	10: return S	Q.	if $\vdash \mathbb{T}_{in} \triangleright \mathbb{T}^{P} \land \vdash \mathbb{T}^{P} \triangleleft \mathbb{T}_{in}$ then
798	(a) Top-level synthesis algorithm.	10:	$S \leftarrow S \cup \{P\}$
799		11.	continue:
800	1: procedure TypeIncompatible(<i>P</i>)	10.	for all $P' \in \text{Expand}(G, P)$ do
801	input: A partial program <i>P</i>	12:	if VIOLATES EMMA $(P' \Phi)$ then
802	output: True if type-incompatible	14.	$\frac{1}{2} \frac{1}{2} \frac{1}$
803	2: for all $n \in Nodes(P)$ do	14:	continue; alas : f Tupply contracts $p(D')$ there
804	3: if IsComplete $(P(n))$ then	15:	eise II TYPEINCOMPATIBLE(P) then
805	4: $\mathbb{T} \leftarrow \text{TypeOf}(P(n))$:	16:	$\Phi \leftarrow \Phi \cup \text{INFERLEMMAS}(P^*, \mathbb{1}_{in});$
806	5: if $\vdash \mathbb{T} \not\sim \text{GoalType}(n)$ then	17:	else
807	6: return true:	18:	$\mathcal{W} \leftarrow \mathcal{W} \cup \{P'\};$
007	return false;	19:	return S ;
000	(b) Procedure for checking program infeasibility .		(c) Goal type synthesis algorithm.
809	51 6 7		··· · •

Fig. 16. Procedures for program synthesis. In SYNTHESIZEVIS, \mathcal{G}_v is the grammar for the plotting sub-DSL and \mathcal{G}_t is the grammar for the table transformation sub-DSL. $\mathbb{T}_t \wedge \mathbb{T}_{in}^p$ stands for the intersection type of \mathbb{T}_t and \mathbb{T}_{in}^p . We provide the procedure for computing type intersection in the appendix.

synthesis. Hence, such lemmas can be useful both for proving the unrealizability of a top-level synthesis goal as well as pruning the search space during synthesis.

Given a lemma (G, R), we refer to G as the *guard* of the lemma and R as the *requirement*. We also say that a lemma is *activated* if the goal type of the synthesis task is a subtype of G. Clearly, the more general the guard of the lemma, the more pruning opportunities that lemma provides.

Pruning with lemmas. To understand how our synthesis procedure utilizes such lemmas, observe that the SYNTHESIZEGOAL algorithm from Figure 16c invokes the VIOLATESLEMMA procedure shown in Figure 17a. Given a partial program P and a lemma $(G, R) \in \Phi$, this procedure checks if there exists some hole in P that violates that lemma. In particular, a hole h with annotated goal type \mathbb{T} violates the lemma if G is activated (i.e., $\mathbb{T} <: G$) and \mathbb{T} is incompatible with requirement R. If this type compatibility check fails for *any* of the holes, then P guaranteed to be infeasible.

THEOREM 2. Let P be a partial program with input type \mathbb{T}_{in} for table D and whose top level goal type is \mathbb{T}_{out} . If VIOLATESLEMMA(P, Φ) returns true, then P(D) is not an inhabitant of \mathbb{T}_{out} .

Learning lemmas. Next, we discuss how to use our refinement type system to infer these
 synthesis lemmas. As shown in line 16 of Figure 16c, our synthesis technique invokes a procedure
 called INFERLEMMAS (presented in Figure 17b) every time it encounters an infeasible partial program.
 In order to generate useful lemmas, we introduce the notion of *type interpolants* that are inspired

834 input: A partial program P input: A failed partial program P 835 **input:** A set of lemmas Φ **input:** Input type \mathbb{T}_{in} of P 836 output: true if *P* is infeasible, false otherwise **output:** A set of learned lemmas Φ 837 for all $h \in Holes(P)$ do $\Phi \leftarrow \{\};$ 2: 2: 838 for all $n \in \text{CompleteNodes}(P)$ do for all $(G, R) \in \Phi$ do 3: 3: 839 4: **if** \vdash GoalType(h) <: G **then** 4: $\mathbb{T} \leftarrow \mathsf{TypeOf}(P(n));$ **if** \vdash GoalType(*h*) $\not\sim$ *R* **then if** \vdash GoalType(n) \nsim T **then** 840 5: 5: return true; $G_n \leftarrow \text{GetInterpolant}(n)$ 6: 841 6: 7: $\tau_{in}, \tau_{G_n} \leftarrow \text{GetBaseTypes}(\mathbb{T}_{in}, G_n)$ 7: return false; 842 8: $R_n \leftarrow \text{GenReq}(\tau_{in}, \tau_{G_n}, \text{max_depth});$ 843 (a) Procedure for checking violation of lemmas. 9: $\Phi \leftarrow \Phi \cup (G_n, R_n);$ 844 10: return Φ ; 845 (b) Procedure for inferring lemmas. 846 Fig. 17. Core type-directed lemma learning procedures 847 848 849 by Craig interpolation [Craig 1957] in logic. Intuitively, type interpolants allow our algorithm to 850 learn lemmas with generalizable guards that can be activated in many contexts. 851 **Definition 6.3 (Base type interpolant).** Given two incompatible base types τ_1 and τ_2 , we say that 852 τ is a base type interpolant for τ_1 and τ_2 if (1) $\tau_1 <: \tau$, (2) $\tau \neq \tau_2$, and (3) for any τ' such that $\tau <: \tau'$, 853 we have $\tau' \sim \tau_2$. 854 **Example 6.1.** If τ_1 = Table({colA : Discrete, colB : Qualitative, colC : Continuous}) and τ_2 = 855 856 Table({colA : Qualitative, colB : Qualitative, colC : Continuous}) then the base type interpolant for 857 τ_1 and τ_2 is Table({colA : Quantitative}). Note that the type interpolant isolates the incompatibility; 858 namely colA in τ_1 's schema is a Quantitative data type, but colA in τ_2 's schema is Qualitative. 859 Next, we generalize this notion from base types to refinement types: 860 **Definition 6.4 (Refinement type interpolant).** Given two refinement types $\mathbb{T}_1 = \{v : \tau_1 \mid \phi_1\}$ 861 and $\mathbb{T}_2 = \{ v : \tau_2 \mid \phi_2 \}$, we say that \mathbb{T} is a type interpolant for \mathbb{T}_1 and \mathbb{T}_2 if: 862 • $\tau_1 \neq \tau_2$, then \mathbb{T} is the base type interpolant for τ_1 and τ_2 863 • $\tau_1 \sim \tau_2$, then $\mathbb{T} = \{ v : \tau_1 \mid \phi \}$ and ϕ is a Craig interpolant for ϕ_1 and ϕ_2 864 865 **Example 6.2.** Let $\mathbb{T}_1 = \{v : \text{Table}(\{\text{colA} : \text{Discrete}, \text{colB} : \text{Discrete}\}) \mid |(v, \{\text{colA}\})| \leq |(v, \{\text{colB}\})| \leq |(v, \{\text{colB}\}$ 866 20) and $\mathbb{T}_2 = \{v : \text{Table}(\{\text{colA} : \text{Discrete}\}) \mid |(v, \{\text{colA}\})| = 30\}$. Then $\{v : \text{Table}(\{\text{colA} : v \in \mathbb{C}\}) \mid v \in \mathbb{C}\}$ 867 Discrete, colB : Discrete}) $||(v, \{colA\})| \le 20\}$ is a refinement type interpolant for \mathbb{T}_1 and \mathbb{T}_2 . 868 With these definitions in place, we now describe INFERLEMMA in more detail. Given an infeasible 869 partial program P, INFERLEMMA first iterates over every complete node n in P and checks whether 870 *n*'s goal type and actual type are incompatible (lines 4-5). If they are, it proceeds to generate a 871 lemma (G_n, R_n) where G_n is a type interpolant between n's goal and actual types (lines 6) and the 872 requirement R_n is generated using the call GENREQ (line 8). Intuitively, the use of type interpolants 873 allows learning lemmas whose guards are as general possible so that they are frequently activated. 874 The GENREQ procedure for generating a requirement is presented as inference rules in Figure 18. 875 At a high level, GENREQ infers DSL constructs that must be used in order to satisfy the goal type 876 and expresses these as syntactic constraints. In more detail, this procedure takes three inputs: (1) 877 the base type τ_{in} for input table D, (2) the base type τ_{out} of the lemma guard, and (3) a synthesis 878 depth k which serves as an upper-bound on the AST depth of the program to be synthesized. The 879 output of GENREQ is a refinement type R such that all programs of maximum AST depth k and 880 with base type $\tau_{in} \rightarrow \tau$ where $\tau \sim \tau_{out}$ must have an output type that is compatible with *R*. 881

1:18 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

1: procedure INFERLEMMAS(P, T_{in})

1: **procedure** VIOLATESLEMMA(P, Φ)

$$BASE \frac{k = 1}{k + (\tau_{in}, \tau_{out}) \rightsquigarrow \{v : \tau_{out} \mid \phi_{R1}\}}{k + (\tau_{in}, \tau_{out}) \rightsquigarrow \{v : \tau_{out}, + \tau'_{out} \sim \tau_{out}\}} REC \frac{k = 1}{k + (\tau_{in}, \tau_{out}) \rightsquigarrow \{v : \tau_{out}, + \tau'_{out} \sim \tau_{out}\}}{k + (\tau_{in}, \tau_{out}) \rightsquigarrow \{v : \tau_{out} \mid \psi_{R1}\}} \sqrt{k + (\tau_{in}, \tau_{out}) \rightsquigarrow \{v : \tau_{out}, + \tau'_{out} \sim \tau_{out}\}}}$$

Fig. 18. GENREQ procedure where k describes an upper bound on the maximum AST depth of the function to be synthesized. τ_{in} is the desired input type and τ_{out} is a base type that the output must be compatible with.

We now explain the two inference rules from Figure 18 in more detail. The first rule, labeled BASE, is the base case for the recursive GENREQ procedure. In the case where k = 1, GENREQ finds the set of F of all DSL operators f such that f takes as input a value of base type τ_{in} and produces an output whose base type is compatible with τ_{out} . Then, the generated requirement is that the synthesized function must contain one of the operators in F: this is expressed as a disjunction of syntactic constraints, where each formula is of the form $\bigvee_{c_i} \pi(v.c_i, f)$ and c_i is an index over the attributes of τ_{out} . Intuitively, this formula says that f could be used to derive any of the columns in the target table's schema.

Example 6.3. Suppose τ_{in} = Table({colA : Qualitative}) and τ_{out} = Table({colA : Discrete}). When k = 1, GENREQ returns {v : Table({colA : Discrete}) | $\pi(v.colA, count)$ } as count is the only operation which directly transforms a Qualitative column to a Discrete one.

The second rule from Figure 18 handles the case for k > 1. To compute a suitable requirement, it first utilizes the base case to get an encoding ϕ_{R_1} of all programs of depth 1 whose input type is τ_{in} to whose output is compatible with τ_{out} . Next, it computes an encoding ϕ_{R_2} of all programs of depth $k \ge 2$ of the form $P^{k-1} \circ f^1$ where f^1 is a function from τ_{in} to an intermediate type τ_t , and P^{k-1} is a program of depth at most k - 1 whose input type is τ_t and output type is compatible with τ_{out} . Thus, the constraint $\phi_{R_1} \lor \phi_{R_2}$ encodes the requirement for all programs up to depth k.

Example 6.4. Suppose τ_{in} = Table({colA : Qualitative}) and τ_{out} = Table({colA : Continuous}). Then ϕ_{R_1} is \perp (false) because there is no operation that can directly transform a Qualitative column to a Continuous one. However, a Qualitative column can only be converted to a Continuous one via the count operation followed by a mean or sum. As such, $\phi_{R_2} = \pi(v.colA, count) \land (\pi(v.colA, mean) \lor \pi(v.colA, sum))$. Thus, GENREQ returns {v : Table({colA : Continuous}) | ϕ_{R_2} }

We now state and prove theorems about our main synthesis procedure SYNTHESIZEVIS.

THEOREM 3. (Soundness) Suppose SynthesizeVis($(\mathbb{T}_p, \mathbb{T}_t), D$) returns a set of programs S. Then for each visualization program $P_v = P_p \circ P_t \in S$, $P_t(D) \models \mathbb{T}_t$ and $P_v(D) \models \mathbb{T}_p$.

THEOREM 4. (Completeness) Given a specification $(\mathbb{T}_p, \mathbb{T}_t)$ and input table D, if there is a visualization program $P_v = P_p \circ P_t$ such that $P_t(D) \models \mathbb{T}_t$ and $P_v(D) \models \mathbb{T}_p$, then $P_v \in SYNTHESIZEVIS((\mathbb{T}_p, \mathbb{T}_t), D)$

7 IMPLEMENTATION

We have implemented the proposed algorithm as a new tool called Graphy written in Python. In what follows, we describe key implementation details that are not covered in the technical sections.

Parser implementation and training. Our NL parser is based on the BERT implementation and pre-trained weights provided by HuggingFace [Wolf et al. 2020]. For training, our model is jointly trained on a collection of examples $(N, \{C_i^*, c_i^*\}_i)$ where we observe the goal properties and

values for each natural language utterance. Our training loss for an example is:

$$\mathcal{L}(N, \{C_i^*, c_i^*\}_i) = \sum_i -\log p(C_i^* \mid N, D, i; \mathbf{w}_{c_i}) - \log p(c_i^* \mid N, D, i; \mathbf{W})$$

which is the standard negative log likelihood objective. We can optimize this objective with standard stochastic gradient descent, simultaneously training the shared BERT encoder parameters as well as the weight vectors and matrices comprising the classification layers. To implement the training procedure, we use the AdamW optimizer [Loshchilov and Hutter 2019], and train the models for 20 epochs with a batch size of 16. We provide the detailed hyperparameters including the learning rates in the appendix.

Type interpolants. Recall that our lemma generation technique from Section 6.2 utilizes the notion of type interpolants to generate guards. While computing interpolants for base types is quite straightforward, we sometimes also need to compute Craig interpolants for the logical qualifiers. In order to ensure that the overhead of this procedure does not outweigh its benefits, we use a simple template-based approach to generate interpolants. In particular, we only generate interpolants that are conjunctions of predicates and enumerate them in increasing number of atomic predicates up to a small bound. The atomic predicates are generated from a pre-defined family of predicates and instantiated with terms and constants that appear in the input formulas.

Ranking visualizations. As described in Section 5, Graphy ranks specifications based on the probabilities produced by the intent classifier and slot filling model. To break ties between programs associated with the same specification, we use the order in which they were explored during the synthesis process, which has the effect of ranking simpler programs above more complicated ones.

8 EVALUATION

We now describe a series of experiments designed to answer the following research questions:

- RQ1. How do the results produced by Graphy compare against those of existing tools?
- RQ2. How long does Graphy take to synthesize visualizations?
- RQ3. How important are the refinement type system and lemma learning for performance?
- RQ4. How effective do users find Graphy in generating visualizations?

Benchmarks. To answer these questions, we perform an evaluation on the NLVCORPUS benchmark suite [Srinivasan et al. 2021]. NLVCORPUS contains a large collection of *real-world* natural language queries and their corresponding ground truth visualizations for three domains, namely Cars, Movies, and Superstore. In total, the corpus contains over 700 queries, gathered from around 200 users and specifying a variety of visualizations. We also note that the queries in this benchmark set are quite diverse as they are *syntactically unrestricted*, and each user was only allowed to specify the visualizations from one of the domains. Table 1 gives a high level summary of the NLVCORPUS benchmarks.

Training set for the parser. In order to use our parser, recall that we first need to train it. Hence, to evaluate Graphy on one of the domains (e.g., Cars) of NLVCORPUS, we train it on the other two domains (e.g., Movies and Superstore). The training data for each of the domain is automatically

974			
	Domain	# of Columns	NL Queries
975	Cars	9	278
976	Movies	10	243
977	Superstore	27	209
978			730

Domain	Parsing Time	Synthesis Time	Total Time
Cars	4.81	0.67	5.48
Movies	4.89	0.94	5.83
Superstore	5.82	0.69	6.51
	5.13	0.77	5.89



Table 2. Average Graphy runtime in seconds.

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

7	Teel		Cars		Movies			Superstore		
1001		top-1	top-5	top-10	top-1	top-5	top-10	top-1	top-5	top-10
Pula basad	NL4DV	0.43	0.49	0.49	0.43	0.48	0.49	0.05	0.46	0.51
Kule-Daseu	DRACO-NL	0.36	0.53	0.57	0.26	0.40	0.41	0.40	0.59	0.59
	NCNET-ORIGINAL	0.08	0.08	0.08	0.09	0.09	0.09	0.07	0.07	0.07
Translation-based	NcNet-Augmented	0.10	0.11	0.11	0.12	0.12	0.12	0.07	0.07	0.07
	BART-VIS ³	0.09	0.11	0.12	0	0.08	0.15	0	0	0
Gi	aphy	0.58	0.77	0.85	0.48	0.64	0.71	0.54	0.81	0.84

Table 3. Comparison between Graphy and other tools in terms of accuracy on the NLVCORPUS.

generated from the corresponding ground truth visualization programs provided by NLVCORPUS with no manual effort required.

Experimental Setup. All of our experiments are conducted on a machine with Intel Xeon(R) W-3275 2.50 GHz CPU and 32 GB of physical memory, running the Ubuntu 18.04 operating system with a NVIDIA Quadro RTX8000 GPU.

Graphy Configuration. In all of our experiments, we configure Graphy to terminate after it finds ten visualization programs. We sort these programs by the score obtained from the parser and break ties by prioritizing programs with smaller AST sizes.

8.1 Comparison with Other Tools

1000 To answer our first research question, we compare Graphy against the following existing tools:

- NL4DV [Narechania et al. 2021]: A state-of-the-art *rule-based* technique for generating visualizations from natural language.
- DRACO-NL [Moritz et al. 2019]: A variant of DRACO, which is a visualization recommendation system that generates visualizations from a partial specification. Even though DRACO does not support natural language queries by default, we implemented a custom translator that converts the output of our NL parser to partial specifications in DRACO's query language.
- NCNET-ORIGINAL [Luo et al. 2022]: A transformer based encoder-decoder model which translates natural language queries to visualizations. This variant of NCNET was trained only on the NL2VIS dataset [Luo et al. 2021].
- NCNET-AUGMENTED: A variant of NCNET that was trained on an augmented dataset that combines both NL2VIS and NLVCORPUS. To ensure there is no overlap between the training and test data, we test on one of the domains (e.g., Cars) from NLVCORPUS and train on the other two (e.g., Movies and Superstore) when performing our evaluation.
- BART-VIS [Lewis et al. 2020]: A translation-based approach that uses a fine-tuned BART language model to directly generate visualizations. We adopt a similar training and testing set-up as NCNET-AUGMENTED.

Main results. We evaluate the performance of all tools in terms of their average top-1, top-5 and top-10 accuracy with respect to the ground truth label. As summarized in Table 3, Graphy outperforms all other tools in terms of accuracy. Among these tools, NL4DV and DRACO-NL are the closest competitors to Graphy; however, they both have significantly lower top-10 results compared to Graphy, and their performance fluctuates across different domains.

Running time. As demonstrated in Table 3, Graphy has better overall accuracy across the board; however, the reader may wonder if this accuracy comes at the cost of significantly longer running times. The last column in Table 2 shows the average end-to-end running time of Graphy across

1029

1017

1023

1024

1025

1026

987 988

989

990

991

992 993

994

995

996

997 998

¹⁰²⁷ ³While the number for this baseline is low, we did confirm that BART-VIs gave 31% accuracy on the test set of NL2VIS.
¹⁰²⁸ However it does not seem to train well with the scale of data we have in the real-world setting.

1:22 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig



the three different domains. As we can see from this table, Graphy is quite fast despite performingenumerative synthesis, taking an average 5.89 seconds to complete each benchmark.

Failure analysis for the baselines. As we can see from Table 3, some of the baseline tools perform quite poorly on the NLVCORPUS data set, so we try to provide some intuition about why this is the case. At a high level, machine translation-style approaches do not do well for two main reasons. First, they require the natural language query to have a complete specification of the intended plot; however, many of the queries in NLVCORPUS only have partial specifications, similar to the working example from Section 2. Second, machine translation approaches do not take any logical constraints into account and may therefore end up generating non-sensical visualizations, such as a bar chart where the x-axis is associated with a continuous variable. Among the rule-based techniques, NL4DV heavily relies on the both the parsing and visualization recommendation rules encoded in the tool and therefore fails to achieve good results when the dataset becomes more complicated and rules cannot generalize. Finally, while DRACO-NL utilizes the output of our parser, we observed that it produces low-quality results when the ground truth visualization requires performing non-trivial aggregation operations over the input data set.

Failure analysis for Graphy. We also analyzed the cases where Graphy does not produce the intended visualization among its top-*k* results. In some cases, the ground truth visualization was not ranked sufficiently high, but Graphy can generate the intended visualization as we increase the value of *k*. In most cases, however, Graphy fails to generate the correct visualization because the natural language query does not contain enough hints about the attributes that should be used to generate the visualization. In such cases, the parser is not able to infer even the base type for the output of the table transformation program, resulting in a very imprecise specification.

8.2 Ablation Study

In this section, we present the results of an ablation study to justify some of the design choices underlying Graphy.

8.2.1 Importance of refinement types. First, we quantify the impact that each component of our refinement type system has on Graphy's ability to prune infeasible programs. We do so by comparing the following three variants of Graphy:

- GRAPHY-BASEONLY: This is a variant of Graphy that only uses base types, but no logical qualifiers.
- GRAPHY-SYNONLY: This is a variant of Graphy that uses the base types and the syntactic constraints in the type system, but no table property constraints.
- GRAPHY-TABLEONLY: This is a variant of Graphy that uses the base types and the table property predicates in the type system, but no syntactic constraints.

Given unbounded time, all of these variants will have the same accuracy as Graphy because they all check that a candidate program satisfies the specification by running it on the input table. As such, we compare the time they take to complete each benchmark (i.e., return ten visualization

programs). In particular, we run each variant over all the benchmarks (with a 60 second timeout),and record the number of benchmarks each one completes along with the average time taken.

The results of this ablation study are shown in Table 4, where we report both the average synthesis time in seconds (excluding the time to parse the natural language description into a refinement type) as well as the percentage of benchmarks solved within the 60 second time limit. Compared to BASEONLY, SYNONLY is almost 10× faster and TABLEONLY completes nearly 50 more benchmarks within the time limit. Finally, having both syntactic and table property constraints allows Graphy to complete all the benchmarks (66 more benchmarks than SYNONLY) and provides an overall speedup of 28× compared to base types alone.

8.2.2 Importance of lemma learning. We also perform a second ablation study to evaluate the importance of the type-directed lemma learning technique presented in Section 6.2. To perform this study, we consider a variant of Graphy called GRAPHY-NOLEMMA that is the same as Graphy except that it does not perform type-directed lemma learning.

The results of this ablation study are presented in Figure 19, which shows the number of benchmarks completed (x-axis) within a given time limit (y-axis) when generating top-10 visualizations. As we can see from the gap between the two lines, Graphy is significantly faster than GRAPHY-NOLEMMA and achieves an overall speed of 2.1× across all benchmarks. For example, Graphy can complete 97% of the benchmarks within 2 seconds, whereas GRAPHY-NOLEMMA completes 77%.

8.3 User Study

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097 1098

1113

1114

1115

1116

1117

1118

We conducted a small user study to evaluate whether Graphy is helpful to end users. We recruited 12 participants, consisting of a mix of undergraduate and graduate students in computer science, math, and business. We asked each participant to reproduce 2 plots ⁴ from the Cars domain using both Graphy and Excel. For each plot and tool, we gave the participants 15 minutes to reproduce the plot using the tool.

Graphy Setup. To facilitate this user study, we developed a UI on top of Graphy. The UI allows
 users to enter a natural language query and presents the top 10 visualizations generated by Graphy.
 To avoid biasing the type of natural language query, no examples were given; participants were
 only told to keep the query high-level and under twenty words.

Excel Setup. We gave each participant a 10 minute tutorial demonstrating how to generate a scatter plot and perform table transformations using PivotTable. Participants were also allowed to use any online resource of their choice.

Results. When using Excel, participants could only finish 92% of the tasks within the time limit, and took 443 seconds on average to solve a task. On the other hand, when using Graphy, the participants could solve *all* the tasks and were, on average, nearly 12.7× faster. We give a more thorough presentation of our user study in the appendix.

9 RELATED WORK

NLIs for data visualization. Many visualization NLIs are powered by (1) ruled-based translation
 engines [Gao et al. 2015; Sun et al. 2010; Yu and Silva 2020b] that pattern-match keywords in the
 user input and translate them into visualization constructs, or (2) neural translation engines that
 leverage encoder-decoder models [Luo et al. 2021, 2022] or pre-trained language models [Poesia
 et al. 2022] to directly generate a visualization program. Because existing systems expect the input
 NL to be complete specifications of the visualization task, they do not perform well in complex

¹¹²⁶ ⁴the plots were of different styles and required different aggregation operations to derive them.

1:24 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

tasks where the query is incomplete or the task requires data transformation [Narechania et al.
2021]. Graphy addresses this issue by formulating the visualization task as a recommendation task:
it first extracts an incomplete user specification from the NL query and then generates diverse
recommendations from it. As shown in our evaluation, Graphy generalizes better to complex tasks
and improves user experience.

1134 Visualization recommendation systems. Visualization recommendation systems are built to help 1135 the user explore the visualization design space from incomplete specifications. For example, 1136 DRACO [Moritz et al. 2019] leverages a constraint solver to recommend visualizations based on the 1137 user's design and data constraints written in answer set programs; Voyager [Wongsuphasawat 1138 et al. 2015] and ShowMe [Mackinlay et al. 2007] use heuristics to recommend visualization chart 1139 type and axes based on the user's fields of interests and the data statistics; DeepEye [Qin et al. 1140 2018] is similar to Voyager but with a statistical learning-to-rank model to the rank visualizations. 1141 Unlike existing systems that require formal specifications (e.g., constraints, concrete fields) as input, 1142 Graphy supports visualization recommendation from natural language.

1144 Type-directed program synthesis. Since refinement types [Martin-Lof et al. 1984; Rondon et al. 1145 2008] were introduced, there have been a number of proposed techniques for synthesizing programs 1146 from refinement type specifications [Frankle et al. 2016; Knoth et al. 2019; Knowles and Flanagan 1147 2009; Osera 2019; Osera and Zdancewic 2015; Polikarpova et al. 2016]. In particular, MYTH2 takes a 1148 function type, along with input-output examples and generates a refinement type specification 1149 by combining the type signature and examples. SyNOUD [Polikarpova et al. 2016] synthesizes 1150 programs using polymorphic refinement types. At the heart of their procedures is a round-trip type 1151 checking mechanism that interleaves top-down and bottom-up propagation of type information. 1152 Graphy employs a similar approach but with two key differences: First our approach propagates 1153 necessary conditions to ensure *type-compatibility* as opposed to subtyping, and second, we apply 1154 type-directed lemma learning to further speedup synthesis over multiple specifications. 1155

1156 Program synthesis from NL. Beyond data visualization, there have also been proposals for per-1157 forming program synthesis directly from natural language [Brown et al. 2020; Lin et al. 2020; 1158 Yaghmazadeh et al. 2017; Ye et al. 2021]. These techniques can mainly divided into two categories: 1159 end-to-end parsing vs parse-then-synthesize techniques. Most of the recent work from the NLP 1160 community focuses on end-to-end parsing, using either powerful generic language models [Brown 1161 et al. 2020; Lewis et al. 2020] or domain-specific techniques targeting SQL [Lin et al. 2020; Wang 1162 et al. 2020], spreadsheet formulas [Gulwani and Marron 2014], and bash commands [Lin et al. 1163 2018]. Parse-then-synthesize approaches, of which Graphy is an instance, first parse the natural 1164 language into an intermediate specification such as a sketch [Yaghmazadeh et al. 2017] or a func-1165 tion declaration [Gvero and Kuncak 2015] and then synthesize programs from these intermediate 1166 specifications. 1167

Lemma learning. Graphy's type directed lemma learning strategy is most similar to the conflict analysis procedures in CDCL-based synthesizers such as Neo and Concord [Chen et al. 2020; Feng et al. 2018]. In particular, these conflict analysis procedures generate a conflict clause whenever the synthesizer determines a partial program is infeasible, and this clause is used prune many other infeasible partial programs. However, unlike Graphy's lemmas, these conflict clauses are only useful for a single synthesis task whereas Graphy's lemmas can be reused across *multiple* synthesis tasks so long as they use the same input dataset.

1176

1168

1133

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

1177 10 CONCLUSION

We have presented Graphy, a new synthesis-based NLI for visualizations. We evaluated Graphy
 on 3 datasets with over 700 natural language queries and found it significantly outperforms prior
 state-of-the-art approaches in top-1, top-5, and top-10 accuracy.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. nnnnnn and Grant No. mmmmmmm. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

1189 REFERENCES

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations*. https://arxiv.org/abs/1409.0473
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan,
 Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan,
- Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz
 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and
- Dario Amodei. 2020. Language Models are Few-Shot Learners. In Advances in Neural Information Processing Systems,
 H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf
- Yanju Chen, Chenglong Wang, Osbert Bastani, Isil Dillig, and Yu Feng. 2020. Program Synthesis Using Deduction-Guided Reinforcement Learning. In *Computer Aided Verification: 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21–24, 2020, Proceedings, Part II* (Los Angeles, CA, USA). Springer-Verlag, Berlin, Heidelberg, 587–610. https: //doi.org/10.1007/978-3-030-53291-8_30
- William Craig. 1957. Linear reasoning. A new form of the Herbrand-Gentzen theorem. Journal of Symbolic Logic 22, 3 (1957), 250–268. https://doi.org/10.2307/2963593
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao,
 Alexander Rudnicky, and Elizabeth Shriberg. 1994. Expanding the Scope of the ATIS Task: The ATIS-3 Corpus. In
 Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994. https:
 //aclanthology.org/H94-1010
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423
- Yu Feng, Ruben Martins, Osbert Bastani, and Isil Dillig. 2018. Program Synthesis Using Conflict-Driven Learning. 53, 4 (jun 2018), 420–435. https://doi.org/10.1145/3296979.3192382
- John K. Feser, Swarat Chaudhuri, and Isil Dillig. 2015. Synthesizing Data Structure Transformations from Input-Output Examples. *SIGPLAN Not.* 50, 6 (jun 2015), 229–239. https://doi.org/10.1145/2813885.2737977
- 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
- Tong Gao, Mira Dontcheva, Eytan Adar, Zhicheng Liu, and Karrie G. Karahalios. 2015. DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Computing Machinery, New York, NY, USA,* 489–500. https://doi.org/10.1145/2807442.2807478
- Sumit Gulwani and Mark Marron. 2014. NLyze: Interactive Programming by Natural Language for Spreadsheet Data
 Analysis and Manipulation. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data* (Snowbird, Utah, USA) (*SIGMOD '14*). Association for Computing Machinery, New York, NY, USA, 803–814. https://doi.org/10.1145/2588555.2612177
- Sumit Gulwani and Madan Musuvathi. 2008. Cover Algorithms and Their Combination. In Proceedings of the Theory
 and Practice of Software, 17th European Conference on Programming Languages and Systems (Budapest, Hungary)
 (ESOP'08/ETAPS'08). Springer-Verlag, Berlin, Heidelberg, 193–207.
- 1225

1:26 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

- Tihomir Gvero and Viktor Kuncak. 2015. Synthesizing Java Expressions from Free-Form Queries. In *Proceedings of the* 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications
 (Pittsburgh, PA, USA) (OOPSLA 2015). Association for Computing Machinery, New York, NY, USA, 416–432. https://doi.org/10.1145/2814270.2814295
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS Spoken Language Systems Pilot Corpus.
 In Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27,1990. https:
 //aclanthology.org/H90-1021
- Minwoo Jeong and Gary Geunbae Lee. 2006. Exploiting Non-Local Features for Spoken Language Understanding. In
 Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions. Association for Computational Linguistics, Sydney,
 Australia, 412–419. https://aclanthology.org/P06-2054
- 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
- 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
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 7871–7880. https://doi.org/10.18653/v1/2020.acl-main.703
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2020. Bridging Textual and Tabular Data for Cross-Domain Text-to-SQL
 Semantic Parsing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 4870–4888. https://doi.org/10.18653/v1/2020.findings-emnlp.438
- Xi Victoria Lin, Chenglong Wang, Luke Zettlemoyer, and Michael D. Ernst. 2018. NL2Bash: A Corpus and Semantic Parser
 for Natural Language Interface to the Linux Operating System. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA), Miyazaki, Japan.
 https://aclanthology.org/L18-1491
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In International Conference on Learning
 Representations. https://openreview.net/forum?id=Bkg6RiCqY7
- Yuyu Luo, Nan Tang, Guoliang Li, Chengliang Chai, Wenbo Li, and Xuedi Qin. 2021. Synthesizing Natural Language to
 Visualization (NL2VIS) Benchmarks from NL2SQL Benchmarks. Association for Computing Machinery, New York, NY,
 USA, 1235–1247. https://doi.org/10.1145/3448016.3457261
- Y. Luo, N. Tang, G. Li, J. Tang, C. Chai, and X. Qin. 2022. Natural Language to Visualization by Neural Machine Translation.
 IEEE Transactions on Visualization and Computer Graphics 28, 01 (jan 2022), 217–226. https://doi.org/10.1109/TVCG.2021.
 3114848
- 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
- P. Martin-Lof, Z. A. Lozinski, Michael Francis Atiyah, Cecil Arthur Hoare, and J. C. Shepherdson. 1984. Constructive mathematics and computer programming. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 312, 1522 (1984), 501–518. https://doi.org/10.1098/rsta.1984.0073
- 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
- Arpit Narechania, Arjun Srinivasan, and John Stasko. 2021. NL4DV: A Toolkit for Generating Analytic Specifications for
 Data Visualization from Natural Language Queries. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb 2021), 369–379. https://doi.org/10.1109/tvcg.2020.3030378
- 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
- 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
- Guoqiang Pan and Moshe Y. Vardi. 2004. Symbolic Decision Procedures for QBF. In *Proceedings of the 10th International Conference on Principles and Practice of Constraint Programming* (Toronto, Canada) (*CP'04*). Springer-Verlag, Berlin,
 Heidelberg, 453–467. https://doi.org/10.1007/978-3-540-30201-8_34
- Benjamin C. Pierce. 2002. Types and Programming Languages (1st ed.). The MIT Press.
- 1274
- Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

- Gabriel Poesia, Alex Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani. 2022. Syn chromesh: Reliable Code Generation from Pre-trained Language Models. In *International Conference on Learning Representations*. https://openreview.net/forum?id=KmtVD97J43e
- Nadia Polikarpova, Ivan Kuraj, and Armando Solar-Lezama. 2016. Program Synthesis from Polymorphic Refinement Types.
 SIGPLAN Not. 51, 6 (jun 2016), 522–538. https://doi.org/10.1145/2980983.2908093
- Xuedi Qin, Yuyu Luo, Nan Tang, and Guoliang Li. 2018. DeepEye: An automatic big data visualization framework. *Big Data Mining and Analytics* 1, 1 (2018), 75–82. https://doi.org/10.26599/BDMA.2018.9020007
- 1281
 Patrick M. Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid Types. SIGPLAN Not. 43, 6 (jun 2008), 159–169.

 1282
 https://doi.org/10.1145/1379022.1375602
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator
 Networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long
 Papers). Association for Computational Linguistics, Vancouver, Canada, 1073–1083. https://doi.org/10.18653/v1/P17-1099
- Arjun Srinivasan, Nikhila Nyapathy, Bongshin Lee, Steven M. Drucker, and John Stasko. 2021. Collecting and Characterizing
 Natural Language Utterances for Specifying Data Visualizations. 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 464, 10 pages. https://doi.org/10.1145/3411764.3445400
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple
 Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 56 (2014), 1929–1958.
 http://jmlr.org/papers/v15/srivastava14a.html
- Yiwen Sun, Jason Leigh, Andrew E. Johnson, and Sangyoon Lee. 2010. Articulate: A Semi-automated Model for Translating
 Natural Language Queries into Meaningful Visualizations. In Smart Graphics, 10th International Symposium on Smart
 Graphics, Banff, Canada, June 24-26, 2010, Proceedings (Lecture Notes in Computer Science, Vol. 6133), Robyn Taylor, Pierre
 Boulanger, Antonio Krüger, and Patrick Olivier (Eds.). Springer, 184–195. https://doi.org/10.1007/978-3-642-13544-6_18
- Gokhan Tur, Dilek Hakkani-Tür, and Larry Heck. 2010. What is left to be understood in ATIS?. In 2010 IEEE Spoken Language
 Technology Workshop. 19–24. https://doi.org/10.1109/SLT.2010.5700816
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia
 Polosukhin. 2017. Attention is All you Need. In Advances in Neural Information Processing Systems, I. Guyon, U. V.
 Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
 https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020. RAT-SQL: Relation-Aware
 Schema Encoding and Linking for Text-to-SQL Parsers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 7567–7578. https://doi.org/10.18653/v1/
 2020.acl-main.677
- 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
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault,
 Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu,
 Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers:
 State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. https://doi.org/
 10.18653/v1/2020.emnlp-demos.6
- Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2015. Voyager:
 Exploratory analysis via faceted browsing of visualization recommendations. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 649–658.
- Navid Yaghmazadeh, Yuepeng Wang, Isil Dillig, and Thomas Dillig. 2017. SQLizer: Query Synthesis from Natural Language. Proc. ACM Program. Lang. 1, OOPSLA, Article 63 (oct 2017), 26 pages. https://doi.org/10.1145/3133887
- Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2021. Optimal Neural Program Synthesis from Multimodal Specifications.
 In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics,
 Punta Cana, Dominican Republic, 1691–1704. https://doi.org/10.18653/v1/2021.findings-emnlp.146
- Bowen Yu and Claudio T. Silva. 2020a. FlowSense: A Natural Language Interface for Visual Data Exploration within a Dataflow System. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 1–11. https://doi.org/10. 1109/tvcg.2019.2934668
 Bowen Yu and Claudio T. Silva. 2020a. FlowSense: A Natural Language Interface for Visual Data Exploration within a Dataflow System. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 1–11. https://doi.org/10.
- Bowen Yu and Cláudio T. Silva. 2020b. FlowSense: A Natural Language Interface for Visual Data Exploration within a
 Dataflow System. *IEEE Trans. Vis. Comput. Graph.* 26, 1 (2020), 1–11. https://doi.org/10.1109/TVCG.2019.2934668
- 1320John M. Zelle and Raymond J. Mooney. 1996. Learning to Parse Database Queries using Inductive Logic Programming. In1321AAAI/IAAI. AAAI Press/MIT Press, Portland, OR, 1050–1055. http://www.cs.utexas.edu/users/ai-lab?zelle:aaai96
- 1322 1323

1:28	Qiaochu Chen, Shankara Pailoor,	Celeste Barnaby, Abby	Criswell, Chenglong V	Vang, Greg Durrett, and Isil Dillig
------	---------------------------------	-----------------------	-----------------------	-------------------------------------

1324 1325 1326	Luke S. Zettlemoyer and Michael Collins. 2005. Learning to Map Sentences to Logical Form: Structured Classification with Probabilistic Categorial Grammars. In <i>Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence</i> (Edinburgh, Scotland) (UAI'05). AUAI Press, Arlington, Virginia, USA, 658–666.
1327	
1328	
1320	
1329	
1330	
1332	
1333	
1334	
1335	
1336	
1337	
1338	
1339	
1340	
1341	
1342	
1343	
1344	
1345	
1346	
1347	
1348	
1349	
1350	
1351	
1352	
1353	
1354	
1355	
1356	
1357	
1350	
1360	
1361	
1362	
1363	
1364	
1365	
1366	
1367	
1368	
1369	
1370	
1371	
1372	

1373 A PROOF

1383 1384

1398

1399 1400

1401

1402 1403

1404

1405

1406

1407

1408

1412

1413

THEOREM 1. (Soundness of TYPEINCOMPATIBLE) Let P be a partial program with input type \mathbb{T}_{in} and top level goal-type \mathbb{T}_{out} . If TYPEINCOMPATIBLE(P) returns true, then for any completion P' of P with type $\mathbb{T}'_{in} \to \mathbb{T}'_{out}, \vdash \mathbb{T}'_{in} \to \mathbb{T}'_{out} \not\sim \mathbb{T}_{in} \to \mathbb{T}_{out}$.

PROOF. Suppose TYPEINCOMPATIBLE(*P*) returns true. Then by lines 3-5 of Figure 16b, there exists a node $n \in Nodes(P)$ such that the type of the complete program P(n) is incompatible with Goal(n). Since *P'* is a completion of *P*, Nodes(*P*) \subseteq Nodes(*P'*). Hence $n \in Nodes(P')$. Given the subprogram rooted at *n* does not satisfy its goal type, we know that *P'* does not satisfy its top level goal type \mathbb{T}_{out} , i.e. $\nvDash \mathbb{T}'_{out} \sim \mathbb{T}_{out}$. Therefore, $\vdash \mathbb{T}'_{in} \rightarrow \mathbb{T}'_{out} \not \sim \mathbb{T}_{out}$.

THEOREM 2. (Soundness of VIOLATESLEMMA) Let P be a partial program whose top-level goal type is \mathbb{T}_{out} , D an input table, and Φ a set of learned lemmas. If VIOLATESLEMMA(P, Φ) returns true, then for any completion P' of P, P'(D) is not an inhabitant of \mathbb{T}_{out} .

¹³⁸⁸ PROOF. Suppose VIOLATESLEMMA(P,Φ) returns true. Then by lines 2-5 of Figure 17a, there exists ¹³⁸⁹ some hole $h \in \text{Holes}(P)$ and some $(G, R) \in \Phi$ such that $\vdash \text{GoalType}(h) <: G \land \vdash \text{GoalType}(h) \neq R$. ¹³⁹⁰ Given a completion P' of P, let node $n \in \text{Nodes}(P')$ be the node instantiated from h in P with a ¹³⁹¹ terminal symbol. We note the goal type of n to be GoalType(n). Since node n is instantiated from ¹³⁹² h, GoalType(n) = GoalType(h).

Given \vdash GoalType(h) \lt : G, GoalType(h) \nleftrightarrow R, following the definition of the synthesis lemma, we know that \nvDash P'(n)(D): GoalType(h). Then P'(n) is not an inhabitant of its goal type. Since P'(n) is a subprogram of P', this means that P' is not an inhabitant of its top level goal type \mathbb{T}_{out} .

THEOREM 3. (Soundness of SynthesizeVis) Suppose SynthesizeVis($(\mathbb{T}_p, \mathbb{T}_t), D$) returns a set of programs S. Then for each visualization program $P_v = P_p \circ P_t \in S$, $P_t(D) \models \mathbb{T}_t$ and $P_v(D) \models \mathbb{T}_p$.

PROOF. It follows from line 8 of Figure 16a that a program $P_p \circ P_t$ is only appended to *R* if $P_t(D) \models \mathbb{T}_t$ and $P_p(P_t(D)) \models \mathbb{T}_p$.

LEMMA 4. Let $\mathcal{G}, \mathbb{T}_{in}, \mathbb{T}_{out}, \triangleright, \triangleleft$ be inputs to SYNTHESIZEGOAL, and let D be an input table with $D \models \mathbb{T}_{in}$. Let P be a program such that there exists a complete program P' with most precise type $\mathbb{T}'_{in} \rightarrow \mathbb{T}'_{out}$ that can be derived from P with $P'(D) \models \mathbb{T}_{out}$. Then SYNTHESIZEGOAL will add P to the worklist W.

PROOF. By induction on the number of terminals *m* in the AST of program *P*.

Base Case: m = 0. The only such program P_0 with 0 terminals is a partial program with one hole that is annotated with the goal output type \mathbb{T}_{out} . This program is added to W on line 3 of Figure 16c.

Inductive Hypothesis: Assume this lemma holds for all programs whose ASTs have less than *m* terminals, where $m \ge 0$.

1414 **Inductive Case:** Suppose P_{m+1} has m+1 terminals. Then there is some program $P_{m'}$ with $m' \le m$ 1415 terminals and some production α such that expanding $P_{m'}$ with α produces P_{m+1} .

1416 Since P' can be derived from P_{m+1} by $P_{m+1} \stackrel{*}{\Rightarrow} P'$, P' can also be derived from $P_{m'}$ by $P_{m'} \stackrel{\alpha}{\Rightarrow}$ 1417 $P_{m+1} \stackrel{*}{\Rightarrow} P'$. Thus, by inductive hypothesis, $P_{m'}$ is added to W. Then at some point $P_{m'}$ will be 1419 dequeued from W on line 5 of Figure 16c. The EXPAND procedure on line 8 will identify α as a 1420 possible production, and will expand $P_{m'}$ to P_{m+1} .

1421

1422 Note that $P'(D) \models \mathbb{T}_{out}$ and $D \models \mathbb{T}_{in}$ imply $\vdash \mathbb{T}'_{in} \sim \mathbb{T}_{in} \land \vdash \mathbb{T}'_{out} \sim \mathbb{T}_{out}$, meaning there exists a 1423 completion of P_m such that $\vdash \mathbb{T}'_{in} \to \mathbb{T}'_{out} \sim (\mathbb{T}_{in} \to \mathbb{T}_{out})$. Then, by contrapositive of Theorem 1, 1424 TYPEINCOMPATIBLE(P_{m+1}) will return false.

Similarly, since there exists a completion of P_{m+1} that inhabits its goal type, by contrapositive of Theorem 2, VIOLATESLEMMA(P_{m+1} , Φ) will return false. P_{m+1} will thus be added to W on line 16.

1429 LEMMA 5. (Completeness of SYNTHESIZEGOAL) Let $\mathcal{G}, \mathbb{T}_{in}, \mathbb{T}_{out}, \triangleright, \triangleleft$ be inputs to SYNTHESIZEGOAL, 1430 let D be an input table with type $D \models \mathbb{T}_{in}$, and let \mathcal{S} be the set of programs with respect to the typing 1431 environment Γ returned by SYNTHESIZEGOAL. Then for any complete program P with the most precise 1432 type $\mathbb{T}'_{in} \to \mathbb{T}'_{out}$ in the grammar \mathcal{G} such that $P(D) \models \mathbb{T}_{out}$, $P \in \mathcal{S}$.

PROOF. By Lemma 4, *P* is added to *W*. Note that the only termination condition for the while loop on line 4 of Figure 16c is that we exhaust *W*. Thus, *P* will be dequeued on line 5 of Figure 16c at some point. Since *P* is complete, we check if $\Gamma \vdash \mathbb{T}_{in} \triangleright \mathbb{T}'_{in} \land \Gamma \vdash \mathbb{T}'_{out} \triangleleft \mathbb{T}_{out}$ (line 7). Since *P*(*D*) $\models \mathbb{T}_{out}$, we have $\Gamma \vdash \mathbb{T}'_{out} \triangleleft : \mathbb{T}_{out}$ and $\Gamma \vdash \mathbb{T}_{in} \triangleleft : \mathbb{T}_{in'}$. It then follows that $\Gamma \vdash \mathbb{T}'_{out} \triangleleft : \mathbb{T}_{out}$ and $\Gamma \vdash \mathbb{T}_{in} \triangleleft : \mathbb{T}'_{in}$, so *P* will be added to *S* on line 8 of Figure 16c.

THEOREM 6. (Completeness of SYNTHESIZEVIS) Given a specification $(\mathbb{T}_p, \mathbb{T}_t)$ and input table D, if there is a visualization program $P_p \circ P_t$ such that $P_t(D) \models \mathbb{T}_t$ and $P_p(P_t(D) \models \mathbb{T}_p$, then $P_p \circ P_t \in SYNTHESIZEVIS((\mathbb{T}_p, \mathbb{T}_t), D)$

1443 PROOF. Let P_p and P_t be programs such that $P_t(D) \models \mathbb{T}_t$ and $P_p(P_t(D)) \models \mathbb{T}_p$. From Lemma 5, 1444 we know that P_p is in the set of programs returned by SYNTHESIZEGOAL on line 3 of Figure 16a. 1445 Also by Lemma 5, we know that P_t is in the set of programs returned by SYNTHESIZEGOAL on 1446 line 6 of Figure 16a. Since $P_t(D) \models \mathbb{T}_t$ and $P_v(D) \models \mathbb{T}_p$, it follows from line 9 of Figure 16a that 1447 $P_p \circ P_t \in \text{SYNTHESIZEVIS}((\mathbb{T}_p, \mathbb{T}_t), D).$

1449 B COMPLETE TYPING RULES

¹⁴⁵⁰ B.1 Intersection of Types

Figure 20 presents inference rules that describe how we computing the intersection type of two
 refinement types.

$$\begin{aligned} \text{Subtype} & \frac{\vdash \tau_1 <: \tau_2}{\vdash \tau_1 \land \tau_2 : \tau_1} & \text{Symmetry} & \frac{\vdash \tau_1 \land \tau_2 : \tau'}{\vdash \tau_2 \land \tau_1 : \tau'} \\ & \tau_1 = \text{Table}(\sigma_1) \quad \tau_2 = \text{Table}(\sigma_2) \\ & \sigma_{\text{shared}} = \{c : \tau_1 \land \tau_2 \mid c : \tau_1 \in \sigma_1, \ c : \tau_2 \in \sigma_2\} \\ & \text{Table} & \frac{\tau = \text{Table}(\sigma_{\text{shared}} \cup (\sigma_1 \Delta \sigma_2))}{\vdash \tau_1 \land \tau_2 : \tau} \end{aligned}$$

$$\begin{aligned} \text{Refinement} & \frac{\vdash \tau_1 \land \tau_2 : \tau'}{\Gamma \vdash \{v : \tau_1 \mid \phi_1\} \land \{v : \tau_2 \mid \phi_2\}) : \{v : \tau' \mid \phi_1 \land \phi_2\}} \end{aligned}$$

Fig. 20. Intersection of base and refinement types. $\sigma_1 \Delta \sigma_2$ is the "symmetric difference" between two schemas i.e., the union of different columns

1469 1470

1468

1465 1466 1467

1428

1433

1448

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

1471 $\Gamma \vdash \mathbb{T}_1 \lt: \mathbb{T}_2$ SUB $\frac{\Gamma \vdash e : \mathbb{T}_1}{\Gamma \vdash e \cdot \mathbb{T}}$ 1472 1473 1474 1475 $\Gamma(T) = \{ v : \tau_T \mid \phi_T \}$ 1476 $\Gamma \vdash \tau_T$: Table({ $c_x : \tau_x, c_y$: Quantitative, $c_{color} : \tau_{color}, c_{subplot} : \tau_{subplot}$ }) 1477 $\tau_x \neq \text{Continuous}$ $\tau_{\text{color}} \neq \text{Continuous}$ $\tau_{\text{subplot}} \neq \text{Continuous}$ 1478 $\mathsf{Encode}(\Gamma) \land \mathsf{Encode}(\phi_T) \Rightarrow |(v, \{c_x, c_{\texttt{color}}, c_{\texttt{subplot}}\})| \ge |(v, \{c_y\})|)$ 1479 Bar $\Gamma \vdash Bar(T, c_x, c_y, c_{color}, c_{subplot}) : \{v : BarPlot \mid v : BarP$ $\pi(v.e, T.c_e)$ 1480 1481 $e \in \{x, y, color, subplot\}$ 1482 $\Gamma(T) = \{ v : \tau_T \mid \phi_T \}$ 1483 $\Gamma \vdash \tau_T : \mathsf{Table}(\{c_x : \tau_x, c_y : \tau_y, c_{\mathsf{color}} : \top, c_{\mathsf{subplot}} : \tau_{\mathsf{subplot}}\})$ 1484 $\tau_x \neq \text{Nominal} \quad \tau_y \neq \text{Temporal} \quad \tau_y \neq \text{Nominal} \quad \tau_{\text{subplot}} \neq \text{Continuous}$ 1485 SCATTER 1486 $\Gamma \vdash \text{Scatter}(T, c_x, c_y, c_{\text{color}}, c_{\text{subplot}}) : \{v : \text{ScatterPlot} \mid$ $\pi(v.e, T.c_e)$ 1487 $e \in \{x, y, color, subplot\}$ 1488 1489 $\Gamma(T) = \{ v : \tau_T \mid \phi_T \}$ 1490 $\Gamma \vdash \tau_T : \mathsf{Table}(\{c_\mathsf{x} : \tau_x, c_\mathsf{y} : \mathsf{Quantitative}, c_{\mathsf{color}} : \tau_{color}, c_{\mathsf{subplot}} : \tau_{subplot} \})$ 1491 $\tau_x \neq \text{Nominal} \quad \tau_{\text{color}} \neq \text{Continuous} \quad \tau_{\text{subplot}} \neq \text{Continuous}$ 1492 $\mathsf{Encode}(\Gamma) \land \mathsf{Encode}(\phi_T) \Rightarrow |v, \{c_x, c_{\mathsf{color}}, c_{\mathsf{subplot}}\})| \ge |(v, \{c_y\})|$ 1493 LINE · $\pi(v.e, T.c_e)$ $\Gamma \vdash \text{Line}(T, c_x, c_y, c_{\text{color}}, c_{\text{subplot}}) : \{v : \text{LinePlot} \mid v : v \in v \}$ 1494 $e \in \{x, y, color, subplot\}$ 1495 1496 $\Gamma(T) = \{ v : \tau_T \mid \phi_T \}$ 1497 $\Gamma \vdash \tau_T <: \mathsf{Table}(\{c_{\mathsf{x}} : \tau_x, c_{\mathsf{y}} : \mathsf{Quantitative}, c_{\mathsf{color}} : \tau_{color}, c_{\mathsf{subplot}} : \tau_{subplot} \})$ 1498 $\tau_x \neq \text{Nominal} \quad \tau_{\text{color}} \neq \text{Continuous} \quad \tau_{\text{subplot}} \neq \text{Continuous}$ 1499 $\mathsf{Encode}(\Gamma) \land \mathsf{Encode}(\phi_T) \Rightarrow |\nu, \{c_x, c_{\texttt{color}}, c_{\texttt{subplot}}\})| \ge |(\nu, \{c_y\})|$ 1500 Area 1501 $\Gamma \vdash \text{Area}(T, c_x, c_y, c_{\text{color}}, c_{\text{subplot}}) : \{v : \text{AreaPlot} \mid$ $\pi(v.e.T.c_{e})$ 1502 $e \in \{x, y, color, subplot\}$ 1503 1504 $\Gamma \vdash e : \{v : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{\dots, c_{\text{tgt}} : \tau_{\text{tgt}}, \dots\})$ 1505 $\vdash \tau_{tgt} <: Quantitative \quad \tau' = \tau_t [c_{tgt} \mapsto Discrete]$ 1506 $\phi_1 = \phi \notin \text{Terms}(\phi, c_{\text{tgt}})$ $\phi_2 = \phi_1 \notin \pi(v.c_{\text{tgt}}, \text{bin})$ $\phi' = \phi_2 \wedge |(v, \{c_{tgt}\})| = n \wedge \pi(v.c_{tgt}, bin)$ 1507 Bin — 1508 $\Gamma \vdash \operatorname{bin}(e, n, c_{+\sigma^{+}}) : \{v : \tau' \mid \phi'\}$ 1509 1510 1511 Fig. 21. Typing Rules 1512 1513 Typing Rules **B.2** 1514 Figures 21 and 22 present our complete set of typing rules. 1515 1516 1517

 $\Gamma \vdash e : \{\nu : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{c_0 : \tau_0, \dots, c_n : \tau_n\})$

 $\operatorname{size}(\overline{c_{\operatorname{key}}}) = k \quad \tau' = \operatorname{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k\}), c'_i \in \overline{c_{\operatorname{key}}}$

 $\Gamma \vdash \text{select}(e, \overline{c_{key}}) : \{v : \tau' \mid \phi\}$

 $\Gamma \vdash e : \{v : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{c_0 : \tau_0, \dots, c_n : \tau_n\}), \forall 0 \le i \le n, c_i \ne c_{\text{tgt}}\}$

 $c_{\texttt{tgt}} \notin \overline{c_{\texttt{key}}} \quad \tau' = \texttt{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k, c_{\texttt{tgt}} : \top\}) \quad c'_i \in \overline{c_{\texttt{key}}}$

 $\phi_1 = \phi \notin \text{Terms}(\phi, c_{\text{tgt}})$ $\phi_2 = \phi_1 \notin \pi(v.c_{\text{tgt}}, \text{mutate})$

 $\phi' = \phi_2 \land |(v, \{c_{tgt}\})| \le |(v, \overline{c_{key}})| \land \pi(v.c_{tgt}, mutate)$

 $\Gamma \vdash \text{mutate}(e, c_{\text{tgt}}, op, \overline{c_{\text{kev}}}) : \{v : \tau' \mid \phi'\}$

 $\Gamma \vdash e : \{\nu : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{\dots, c_{\text{tgt}} : \tau_{\text{tgt}}, \dots\})$

 $c_{tgt} \notin \overline{c_{key}}$ $\tau' = \text{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k, c_{tgt} : \tau_{tgt}\})$ $c'_i \in \overline{c_{key}}$ $\vdash \tau_{tgt} : \text{Quantitative} \quad \tau' = \tau' \begin{bmatrix} c_{tgt} \mapsto \text{Continuous} \end{bmatrix}$

 $\phi_1 = \phi \notin \text{Terms}(\phi, c_{\text{tgt}})$ $\phi_2 = \phi_1 \notin \pi(v.c_{\text{tgt}}, \text{mean})$

 $\Gamma \vdash \text{summarize}(e, \overline{c_{\text{kev}}}, \text{mean}, c_{\text{tgt}}) : \{v : \tau' \mid \phi'\}$

 $\Gamma \vdash e : \{v : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{\dots, c_{\text{tgt}} : \tau_{\text{tgt}}, \dots\})$

 $c_{tgt} \notin \overline{c_{key}} \quad \tau' = \text{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k, c_{tgt} : \tau_{tgt}\}) \quad c'_i \in \overline{c_{key}}$ $\tau' = \tau'[c_{tgt} \mapsto \text{Discrete}]$

 $\Gamma \vdash e : \{v : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{\dots, c_{\text{tgt}} : \tau_{\text{tgt}}, \dots\})$

 $c_{\text{tgt}} \notin \overline{c_{\text{key}}} \quad \tau' = \text{Table}(\{c'_0 : \tau'_0, \dots, c'_k : \tau'_k, c_{\text{tgt}} : \tau_{\text{tgt}}\}) \quad c'_i \in \overline{c_{\text{key}}}$

 $\vdash \tau_{tgt} : \text{Quantitative} \quad \tau' = \tau' [c_{tgt} \mapsto \text{Continuous}]$

 $\phi_1 = \phi \notin \text{Terms}(\phi, c_{\text{tgt}})$ $\phi_2 = \phi_1 \notin \pi(v.c_{\text{tgt}}, \text{sum})$

 $\Gamma \vdash \text{summarize}(e, \overline{c_{\text{kev}}}, \text{sum}, c_{\text{tgt}}) : \{v : \tau' \mid \phi'\}$

Fig. 22. Typing Rules

SUMM-MEAN $\frac{\phi' = \phi_2 \wedge |(v, \{c_{tgt}\})| \le |(v, \overline{c_{key}})| \wedge \pi(v.c_{tgt}, \text{mean})}{\Sigma}$

 $\begin{array}{c} \phi_1 = \phi \notin \operatorname{Terms}(\phi, c_{\operatorname{tgt}}) \quad \phi_2 = \phi_1 \notin \pi(v.c_{\operatorname{tgt}}, \operatorname{count}) \\ \phi' = \phi_2 \wedge |(v, \{c_{\operatorname{tgt}}\})| \leq |(v, \overline{c_{\operatorname{key}}})| \wedge \pi(v.c_{\operatorname{tgt}}, \operatorname{count}) \\ \hline \Gamma \vdash \operatorname{summarize}(e, \overline{c_{\operatorname{key}}}, \operatorname{count}, c_{\operatorname{tgt}}) : \{v : \tau' \mid \phi'\} \end{array}$

SUMM-SUM $\frac{\phi' = \phi_2 \wedge |(v, \{c_{tgt}\})| \leq |(v, \overline{c_{key}})| \wedge \pi(v.c_{tgt}, sum)}{\sum |v| = \psi_2 \wedge |(v, \{c_{tgt}\})| \leq |(v, \overline{c_{key}})| \wedge \pi(v.c_{tgt}, sum)}$

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

FILTER $\frac{\phi' = \phi \notin \text{Terms}(\phi, \{c_1, \dots, c_n\})}{\Gamma \vdash \text{filter}(e, val_1 \ op \ val_2) : \{v : \tau \mid \phi'\}}$

 $\Gamma \vdash e : \{v : \tau_t \mid \phi\}$ where $\tau_t = \text{Table}(\{c_0 : \tau_0, \dots, c_n : \tau_n\})$

1521 1522 1523

1520

- 1524 1525
- 1526
- 1527

Select -

1528

1531 1532

1533

MUTATE -

1535

1536

1537
1538

1539

1541

1542 1543

1544 1545 1546

1547

1548 1549

1555

1556

1557

1558 1559

1540

С FORMULA ENCODING 1569

1570 In this section, we describe our Encode procedure which encodes qualifiers in our refinement type 1571 system as formulas in the combined theory of equality, uninterpreted functions, and integers. 1572

Formula Language. Figure 23 presents the syntax of our encoded formulas as a context free 1573 grammar. Note that many of the terminals in our refinement type system also appear in this 1574 grammar but now have different semantics. The symbols $|\cdot|$, Filter and Proj refer to relational 1575 operators in our refinement type system, but correspond to uninterpreted functions in the formula 1576 language. We also import column names c into our formula language, but they refer to object 1577 constants. Finally, in our refinement type system, Proj takes a list of column names as its second 1578 input, whereas in this formula language, Proj takes two inputs where the second argument is an 1579 object constant. 1580

We formalize our encoding procedure as inference rules incorporating judgments of the form.

$$\Sigma \vdash t \rightsquigarrow e, \Sigma'$$

where Σ is an environment that maps terms in our refinement type language to terms in the 1584 formula language. This judgment means: given an environment Σ and a term t in the refinement 1585 type language, Encode returns the corresponding term e in our formula language along with an updated environment Σ' . This formalization is presented in Figure 24. Most of the rules are straightforward, so we highlight the most interesting. 1588

- Syntactic. Our encoding scheme associates each syntactic constraint $\pi(x,\eta,\mu)$ with a unique propositional variable (Syn-1, Syn-2)
- 1591 • **Projection.** For each set of columns \overline{c} , we associate a fresh object constant s. In particular, $\{c_1, c_2\}$ and $\{c_2, c_1\}$ are associated with the same fresh constant as they represent the same 1592 sets (Proj-1, Proj-2). 1593
 - Filter. Every filter operator op is assigned a fresh function constant f_i , and each value val is assigned a fresh object constant s. Thus, the filter operation c op val is treated as the function application $f_i(c, s)$.

Correctness of SYNTHESIZEVIS. Since the combined theory of integers and equality with uninterpreted functions over-approximates the semantics of our qualifiers, one may wonder whether SYNTHESIZEVIS is still complete i.e., doesn't prune correct programs. We now argue that the procedure is still complete. First, we note that our compatibility check will not prune feasible programs since (1) our encoding is the conjunction of formulas that over-approximate their corresponding qualifiers, and (2) our compatibility rule checks that the formula is satisfiable. Second, our subtyping checks are used to prune prune plotting programs whose output type is not a subtype of the goal type specification produced by the parser (line 3 SYNTHESIZEVIS). In that case, the logical qualifiers in the goal type specification are only syntactic constraints, which are *precisely encoded* as boolean constraints. As such, our subtyping check will not prune correct programs there.

1581 1582 1583

1586

1587

1589

1590

1594

1595 1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1618 1619 1620 1621 1622 1623 1624	Formula Language Formula F := Expression E := TableFunction T := ObjectConstants S := FunctionConstants G :=	$\begin{array}{l} \oplus (F_1, \dots, F_n) \mid E \ \emptyset \ E \mid p \\ \mid T \mid \mid Max(T) \mid Min(T) \mid x \mid S \\ Proj(T, S) \mid Filter(T, G(c, S)) \mid x \\ a \mid b \mid \dots \\ f_1 \mid f_2 \mid \dots \end{array}$
1624 1625 1626 1627 1628	Fig. 23. Our formula language where $\oplus \in \{\land, \lor, \neg$ combined theory of equality, uninterpreted funct Min are uninterpreted functions, and <i>p</i> represents	n}, and $\emptyset \in \{=, \le, >\}$. Formulas in this language are in the ions, and integers. In particular, $ \cdot $, Proj, Filter, Max, and s propositional variables.
1629 1630 1631 1632 1633		
1634 1635 1636 1637 1638		
1639 1640 1641 1642 1643		
1644 1645 1646 1647 1648		
1649 1650 1651 1652 1653		
1654 1655 1656 1657		
1659 1660 1661 1662		
1664 1665 1666		

$$\Sigma \vdash \phi_1 \rightsquigarrow F_1, \Sigma_1$$

$$\Sigma_1 \vdash \phi_2 \rightsquigarrow F_2, \Sigma_2$$

Logical Operators
$$\frac{\sum_{n-1} \vdash \phi_n \rightsquigarrow F_n, \sum_n}{\sum_{i} \vdash \oplus(\phi_1, \dots, \phi_n) \rightsquigarrow \oplus(F_1, \dots, F_n), \sum_n}$$

1674
1675
1676
SEMANTIC TERM
$$\frac{\Sigma \vdash \theta_1 \rightsquigarrow E_1, \Sigma_1 \quad \Sigma_1 \vdash \theta_2 \rightsquigarrow E_2, \Sigma_2}{\Sigma \vdash \theta_1 \ \emptyset \ \theta_2 \rightsquigarrow E_1 \ \emptyset \ E_2, \Sigma_2} \quad \text{Card} \quad \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma'}{\Sigma \vdash |\gamma| \rightsquigarrow |T|, \Sigma'}$$
1677

$$Syn-1 \frac{\pi(x.\eta,\mu) \in dom(\Sigma)}{\Sigma \vdash \pi(x.\eta,\mu) \rightsquigarrow \Sigma(\pi(x.\eta,\mu))} \quad Syn-2 \frac{\pi(x.\eta,\mu) \notin dom(\Sigma) \quad \text{fresh } p'}{\Sigma \vdash \pi(x.\eta,\mu) \rightsquigarrow p', \Sigma[\pi(x.\eta,\mu) \leftarrow p']}$$

$$V_{AR-1} \frac{x \in dom(\Sigma)}{\Sigma \vdash x \rightsquigarrow \Sigma(x)} \qquad V_{AR-2} \frac{x \notin dom(\Sigma) \quad \text{fresh } v}{\Sigma \vdash x \rightsquigarrow v, \Sigma[x \leftarrow v]}$$

FILTER OP-1
$$\frac{op \in dom(\Sigma)}{\Sigma \vdash op \rightsquigarrow \Sigma(x)}$$
 FILTER OP-2
$$\frac{op \notin dom(\Sigma), \text{ fresh } f_i}{\Sigma \vdash op \rightsquigarrow f_i, \Sigma[op \leftarrow f_i]}$$

$$\operatorname{Max} \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma'}{\Sigma \vdash \max(\gamma) \rightsquigarrow \operatorname{Max}(T), \Sigma'} \qquad \operatorname{Min} \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma'}{\Sigma \vdash \min(\gamma) \rightsquigarrow \operatorname{Min}(T), \Sigma'}$$

$$\operatorname{ProJ-1} \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma' \quad \overline{c} \notin dom(\Sigma') \quad \operatorname{fresh} s}{\Sigma \vdash \operatorname{Proj}(\gamma, \overline{c}) \rightsquigarrow \operatorname{Proj}(T, s), \Sigma'[\overline{c} \leftarrow s]}$$

$$\operatorname{Proj-2} \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma' \quad \overline{c} \in dom(\Sigma')}{\Sigma \vdash \operatorname{Proj}(\gamma, \overline{c}) \rightsquigarrow \operatorname{Proj}(T, \Sigma'(\overline{c})), \Sigma'}$$

$$\label{eq:Filter-1} \begin{array}{ll} \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma_1 \quad \Sigma_1 \vdash op \rightsquigarrow f_i, \Sigma_2 \quad val \in dom(\Sigma) \\ \vdash \mathsf{Filter}(\gamma, c \ op \ val) \rightsquigarrow \mathsf{Filter}(T, f_i(c, \Sigma(val))), \Sigma_2 \end{array}$$

$$FILTER-2 \frac{\Sigma \vdash \gamma \rightsquigarrow T, \Sigma_1 \quad \Sigma_1 \vdash op \rightsquigarrow f_i, \Sigma_2 \quad val \notin dom(\Sigma) \quad \text{fresh } s}{\vdash \text{Filter}(\gamma, c \ op \ val) \rightsquigarrow \text{Filter}(T, f_i(c, s)), \Sigma_2[val \leftarrow s]}$$

Fig. 24. Inference rules describing the Encode procedure

1:36 Qiaochu Chen, Shankara Pailoor, Celeste Barnaby, Abby Criswell, Chenglong Wang, Greg Durrett, and Isil Dillig

1716 D SEMANTICS OF ½ OPERATOR

¹⁷¹⁷ In this section, we describe the $\frac{1}{2}$ operator introduced in Section 4.4 in more detail. Here we assume ¹⁷¹⁸ the operator takes qualifiers of the form $\phi_s \wedge \phi_p$, where ϕ_s (resp ϕ_p) is a boolean combination ¹⁷¹⁹ of semantic (resp. syntactic) constraints. We give the semantics of $\frac{1}{2}$ as a procedure shown in ¹⁷²⁰ Figure 25. Given a qualifier $\phi_s \wedge \phi_p$ along with a set of terms *T* to remove, we have $\phi_s \wedge \phi_p \frac{1}{2}T \equiv$ ¹⁷²¹ REMOVE $(\phi_s \wedge \phi_p, T)$.

1722 We now describe this procedure in more detail. In line 2 we call Encode (described in Section C) 1723 to encode our qualifiers as logical formulas. Then in lines 4-10, we iterate over all the terms t in 1724 T. If t is a semantic term (line 5), we generate a fresh variable x and replace all occurrences of t 1725 in Φ_s with *x*. We then update Φ_s to be its existential generalization (line 7). On the other hand, if 1726 t is a syntactic term we perform a similar procedure, except Φ_p becomes a QBF formula. Finally, 1727 since the fresh variables we introduced represent the terms we want to forget, we compute the 1728 strongest formulas entailed by ϕ_s (resp. ϕ_p) that don't contain the fresh variables. Since, QBF admits 1729 quantifier elimination [Pan and Vardi 2004], we can derive the strongest QFF formula entailed 1730 by ϕ_p by applying quantifier elimination (line 11). However, since our semantic constraints are 1731 expressed in the combined theory of equality, uninterpreted functions, and integers, which does 1732 not admit quantifier elimination, we instead the Cover algorithm [Gulwani and Musuvathi 2008] 1733 to compute the strongest formula (line 12). Finally, we convert Φ_s and Φ_p back into qualifiers and 1734 return their conjunction (line 13). 1735

Optimization. In our implementation of REMOVE, we apply two optimizations based on the following observations. First, we observe that nearly all our logical qualifiers are conjunctions of literals, and so we represent our qualifiers as sets of literals. As such, when removing a syntactic constraint, we simply remove all corresponding literals from the set. Second, as all our semantic constraints are linear inequalities, we encode our semantic constraints as formulas in Presburger Arithemtic. We then apply Fourier-Motzkin variable elimination when removing semantic terms. We illustrate these optimizations in the examples below.

1743 **Example D.1.** (Forgetting Syntactic Constraints) Suppose we call $\phi_2^{\prime}\pi(v.c_{tgt}, mean)$ where ϕ is 1744 $|v, \{c_1\}| = 30 \land \neg \pi(v.c_{tgt}, mean) \land \pi(v.c_2, count)$. Since this qualifier is a conjunction of literals, 1745 REMOVE maintains a set of constraints $\{|v, c_1| = 30, \neg \pi(v.c_{tgt}, mean), \pi(v.c_2, count)\}$. The only literal 1746 that corresponds to $\pi(v.c_{tgt}, mean)$ is $\neg \pi(v.c_{tgt}, mean)$ and so we drop that literal from the set. Thus, 1747 the formula returned by REMOVE is $|v, c_1| = 30 \land \pi(v.c_2, count)$

Example D.2. (Forgetting Semantic Terms) Suppose we call $\phi_{\pm}|v, \{c_1\}|$ where ϕ is $|v, \{c_1\}| \le |v, \{c_2\}| \land |v, \{c_1\}| = 30$. REMOVE internally constructs an equisat formula $x_1 \le x_2 \land x_1 = 30$ in the theory of integers where x_1 , and x_2 are fresh variables that occur freely and $x_1 \rightarrow |v, \{c_1\}|$, $x_2 \rightarrow |v, \{c_2\}|$. It then applies Fourier-Motzkin variable elimination on x_1 to get the constraint $30 \le x_2$, and then decodes the formula back to the qualifier $30 \le |v, \{c_2\}|$.

- 1754
- 1755 1756
- 1757
- 1758
- 1759
- 1760
- 1761
- 1762
- 1763

- 1765 1: **procedure** REMOVE($\phi_s \land \phi_p, T$)
- ¹⁷⁶⁶ **input:** Logical qualifier of the form $\phi_s \wedge \phi_p$ where ϕ_s is a boolean combination of semantic constraints, ¹⁷⁶⁷ and ϕ_p is a boolean combination of syntactic constraints.
- **input:** A set of terms *T* to forget
- **output:** A qualifier ϕ that does not contain any terms in *T* and whose encoding is the strongest QFF formula that is entailed by the encoding of $\phi_s \wedge \phi_p$.

1771 $\Phi_s \leftarrow \mathsf{Encode}(\phi_s); \Phi_p \leftarrow \mathsf{Encode}(\phi_p)$ 2: $V_{s} = \{\}; V_{p} = \{\}$ 1772 3: for all $t \in T$ do 4: 1773 if lsSemantic(t) then 5: 1774 $x \leftarrow \text{GetFreshVar}()$ 6: 1775 $\Phi_s \leftarrow \exists x. \Phi_s[x/\mathsf{Encode}(t)]; V_s = V_s \cup \{x\}$ 7: 1776 8: else if lsSyntactic(t) then 1777 $p \leftarrow \text{GetFreshPropVar}()$ 9: 1778 $\Phi_p \leftarrow \exists p.\Phi_p[p/\mathsf{Encode}(t)]; V_p = V_p \cup \{p\}$ 10: 1779 $\Phi_p \leftarrow \mathsf{EliminateBoolVars}(\Phi_p, V_p)$ 11: 1780 $\Phi_s \leftarrow \text{ComputeCover}(\Phi_s, V_s)$ 12: 1781 13: **return** $Decode(\Phi_s) \wedge Decode(\Phi_p)$ 1782 1783 Fig. 25. Procedure encoding semantics of $\frac{1}{2}$. 1784

1786 E TRAINING PARAMETERS FOR THE PARSER

Each intent-and-slot-filling model in the parser is trained for 20 epochs using the AdamW optimizer [Loshchilov and Hutter 2019] with a batch size of 16. We use a learning rate of 2e - 05 for BERT, which is one of the standard suggested learning rates [Devlin et al. 2019], a warm-up ratio of 0.05, and an input dropout rate of 0.2 to reduce overfitting [Srivastava et al. 2014]. We trained the models on one NVIDIA Quadro RTX 8000 with 48GB of memory. Each training run of the model took around 10 minutes.

¹⁷⁹⁴ F USER STUDY PROCEDURE

1785

1795

1796 1797

1798 1799

1800

1801

1802

1803

1804

1805

1813

In this section, we describe our user-study protocol in more detail.

User study sessions. Our user study was completed in 12 sessions, one for each participant. The participants used the same laptop, which had Excel and Graphy installed, across all sessions.

Participant Introduction. We started each user study session by first describing the task that the participant needed to accomplish. In particular, we asked them to reproduce two plots shown in Figure 26 using both Excel and Graphy. We chose Excel as the baseline because it is a common data visualization tool that is designed to be accessible to non-expert users. In order to minimize the effect of knowledge transfer, we randomly determined whether a participant was first given access to Graphy or to Excel.

Plot Selection. To avoid biasing the study in Graphy's favor, we randomly selected two plots
 of different types from the Cars domain in NLVCORPUS for the participants to reproduce. To
 ensure that the plots were reasonably challenging we only selected among plots that required data
 aggregation operations in the table transformation.

Dataset Introduction. After instructing the participants on what they needed to do, we showed
 them the relational table CARS that the plots are based on. We gave them 2 minutes to get familiar

Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.





Fig. 26. Two plots used in the user study

with the data set and ask any questions related about it. After they were familiarized, we gave each user a training session for each tool.

Excel Training. We first introduced the Excel spreadsheet interface and showed the participants (1) how to make a scatter plot, and (2) how to do data aggregation using PivotTable, a feature in Excel that enables users to do data aggregation without any coding knowledge. Afterwards, we gave the users 10 minutes to play around with the tool. We encouraged them to try and produce a line chart and a bar chart. In addition to the training session, we also provided a "cheat-sheet" that included Excel documentation that we thought would be helpful to the user when performing the task. Finally, we allowed them to search online for help during the study.

Graphy Training. Like with Excel, we started the training by introducing the participant to Graphy's UI, shown in Figure 27. To avoid biasing the user in any way, we did not present any examples about how to use Graphy but simply asked the user to try it themselves for 5 minutes.

User Study Workflow. Once the participant was familiar with the data set and the tools, we gave them 2 minutes to familiarize themselves with the plots they needed to reproduce. We told the user to let us know when we could start timing and when they thought they had finished the task. In total, each participant had an hour to complete all the tasks using all the tools (15 minutes per task per tool). We made it clear to the participants that they were not required to reproduce *exactly* the same plot as shown in the ground truth, and they could consider themselves to be finished as long as they thought they had produced a plot that conveys the same meaning.

When using Graphy, users would enter a natural language query, and Graphy would return the top-10 results as visualizations back to the user. The user would then skim through the graphs and choose a visualization if they thought was equivalent to the ground truth. If the participant decided none of the visualizations shown was the one they wanted, they could try again by entering a different query. For Excel, we provided the user a spreadsheet that contained the table to be visualized so they did not need to import the data to Excel. During their time working on the plot, the user was allowed to use any resources such as searching the Internet or using the "cheat-sheet" we provide.



1:39