

Automated Migration of Hierarchical Data to Relational Tables using Programming-by-Example

Navid Yaghmazadeh
University of Texas at Austin
nyaghma@cs.utexas.edu

Xinyu Wang
University of Texas at Austin
xwang@cs.utexas.edu

Isil Dillig
University of Texas at Austin
isil@cs.utexas.edu

ABSTRACT

While many applications export data in hierarchical formats like XML and JSON, it is often necessary to convert such hierarchical documents to a relational representation. This paper presents a novel programming-by-example approach, and its implementation in a tool called MITRA, for automatically migrating tree-structured documents to relational tables. We have evaluated the proposed technique using two sets of experiments. In the first experiment, we used MITRA to automate 98 data transformation tasks collected from StackOverflow. Our method can generate the desired program for 94% of these benchmarks with an average synthesis time of 3.8 seconds. In the second experiment, we used MITRA to generate programs that can convert real-world XML and JSON datasets to full-fledged relational databases. Our evaluation shows that MITRA can automate the desired transformation for all datasets.

1. INTRODUCTION

Many applications store and exchange data using a hierarchical format, such as XML or JSON documents. Such tree-structured data models are a natural fit in cases where the underlying data is hierarchical in nature. Furthermore, since XML and JSON documents incorporate both data and meta-data, they are self-describing and portable. For these reasons, hierarchical data formats are popular for exporting data and transferring them between different applications.

Despite the convenience of hierarchical data models, there are many situations that necessitate converting them to a relational format. For example, data stored in an XML document may need to be queried by an existing application that interacts with a relational database. Furthermore, because hierarchical data models are often not well-suited for efficient data extraction, converting them to a relational format is desirable when query performance is important.

In this paper, we introduce a new technique based on *programming-by-example (PBE)* [35, 26] for converting hierarchically structured data to a relational format. In our methodology, the user provides a set of simple input-output examples to illustrate the desired transformation, and our system, MITRA¹, automatically synthesizes a program that performs the desired task. Because MITRA learns the target transformation from small input-output examples, it can achieve automation with little guidance from the user. In a typical usage scenario, a user would “train” our system on a small, but representative subset of the input data and then

use the program generated by MITRA to convert a very large document to the desired relational representation.

While programming-by-example has been an active research topic in recent years [25, 15, 47, 23, 56, 22, 30, 41], most techniques in this space focus on transformations between similarly structured data, such as string-to-string [25, 47], tree-to-tree [56, 23] or table-to-table transformations [22, 51, 30, 41]. Unfortunately, automating transformations from tree- to table-structured data brings new technical challenges that are not addressed by prior techniques. First, because the source and target data representations are quite different, the required transformations are typically more complex than those between similarly-structured data. Second, since each row in the target table corresponds to a relation between nodes in the input tree, the synthesizer needs to discover these “hidden links” between tree nodes.

This paper addresses these challenges by presenting a new program synthesis algorithm that decomposes the synthesis task into two simpler subproblems that aim to learn the column and row construction logic separately:

- *Learning the column extraction logic:* Given an attribute in a relational table, our approach first synthesizes a program to extract tree nodes that correspond to that attribute. In other words, we first ignore relationships between different tree nodes and construct each column separately. Taking the cross product of the extracted columns yields a table that overapproximates the target table (i.e., contains extra tuples).
- *Learning the row extraction logic:* Since the program learned in the first phase produces a table that overapproximates the target relation, the next phase of our algorithm synthesizes a program that filters out “spurious” tuples generated in the first phase. In essence, the second phase of the algorithm discovers the “hidden links” between different nodes in the original tree structure.

Figure 1 shows a schematic illustration of our synthesis algorithm. Given an input tree T and output table R with k columns, our technique first learns k different programs π_1, \dots, π_k , where each *column extraction* program π_i extracts from T the data stored in the i 'th column of R . Our synthesis algorithm then constructs an intermediate table by applying each π_i to the input tree and taking their cross product. Thus, the intermediate table $\pi_1(T) \times \dots \times \pi_k(T)$ generated during the first phase overapproximates the target table (i.e., it may contain more tuples than R). In the next phase, our technique learns a *predicate* ϕ that can be used to filter out exactly the spurious tuples from the intermediate

¹stands for Migrating Information from Trees to RelAtions




Figure 1: Schematic illustration of our approach

table. Hence, the program synthesized by our algorithm is always of the form $\lambda x. \text{filter}(\pi_1 \times \dots \times \pi_k, \phi)$. Furthermore, since the synthesized program should not be over-fitted to the user-provided examples, our method always learns the *simplest* program of this shape that is consistent with the user-provided input-output examples.

From a technical point of view, the contributions of this paper are three-fold. First, we propose a domain-specific language (DSL) that is convenient for expressing transformations between tree-structured and relational data. Our DSL is expressive enough to capture many real-world data transformation tasks, and it also facilitates efficient synthesis by allowing us to decompose the problem into two simpler learning tasks. While the programs in this DSL may sometimes be inefficient, our method eliminates redundancies by memoizing shared computations in the final synthesized program. This strategy allows us to achieve a good trade-off between expressiveness, ease of synthesis, and efficiency of the generated programs.

The second technical contribution of this paper is a technique for automatically learning column extraction programs using *deterministic finite automata (DFA)*. Given an input tree and a column from the output table, our method constructs a DFA whose language corresponds to the set of DSL programs that are consistent with this example. Hence, in our methodology, learning column extraction programs boils down to finding a word (i.e., sequence of DSL operators) that is accepted by the automaton.

The third technical contribution of this paper is a novel technique for learning predicates that can be used to filter out spurious tuples in the intermediate table. Given a set of positive examples (i.e., tuples in the output table) and a set of negative examples (i.e., spurious tuples in the intermediate table), we need to find a smallest classifier in the DSL that can be used to distinguish the positive and negative examples. Our key observation is that this task can be reduced to *integer linear programming*. In particular, our method first generates the universe of all possible atomic predicates over the DSL and then infers (using integer linear programming) a *smallest* subset of predicates that can be used to distinguish the positive and negative examples. Given such a subset, our method then uses standard logic minimization techniques to find a boolean combination of atomic predicates that can serve as a suitable classifier.

We have implemented our technique in a tool called MITRA, which consists of a language-agnostic synthesis core for tree-to-table transformations as well as domain-specific plug-ins. While MITRA can generate code for migrating data from any tree-structured representation to a relational table, it requires plug-ins to translate the input format into our intermediate representation. Our current implementations contains two such plug-ins for XML and JSON docu-

ments. Furthermore, MITRA can be used to transform tree-structured documents to a full-fledged relational database by invoking it once for each table in the target database.

We have evaluated MITRA by performing two sets of experiments. In our first experiment, we use MITRA to automate 98 data transformation tasks collected from Stack-Overflow. MITRA can successfully synthesize the desired program for 94% of these benchmarks, with an average synthesis time of 3.8 seconds. In our second experiment, we use MITRA to migrate four real-world XML and JSON datasets (namely, IMDB, YELP, MONDIAL, and DBLP) to a full-fledged relational database. Our experiments show that MITRA can perform the desired task for all four datasets.

To summarize, this paper makes the following key contributions:

- We propose a programming-by-example approach for migrating tree-structured documents to a relational format.
- We propose a tree-to-table transformation DSL that facilitates synthesis by allowing us to decompose the synthesis task into two subproblems.
- We describe a synthesis technique for learning column transformation programs using deterministic finite automata.
- We present a predicate learning algorithm that reduces the problem to a combination of integer linear programming and logic minimization.
- We implement these ideas in a system called MITRA and empirically demonstrate its practicality and effectiveness.

The rest of the paper is organized as follows: We first provide an overview of our approach through a simple motivating example (Section 2). We then introduce *hierarchical data trees* (Section 3) and our DSL for implementing tree-to-table transformations (Section 4). In Section 5, we present our synthesis algorithm and discuss our implementation in Section 6. Finally, we present our empirical evaluation in Section 7 and survey related work in Section 8.

2. OVERVIEW

In this section, we give a high-level overview of our technique with the aid of a simple motivating example. Consider the XML file from Figure 2a which contains information about the users of a social network as well as the “friendship” relation between them. Suppose a user wants to convert this XML document into the relational table shown in Figure 2b. Observe that this transformation is non-trivial because the XML file stores this information as a mapping from each user to a list of friends, where each friend is represented by their `fid`. In contrast, the desired table stores this information as tuples (A, B, n) , indicating that person with name A is friends with user with name B for n years.

Suppose that the original XML file is much bigger than the one shown in Figure 2a, so the user decides to automate this task by providing the input-output example from Figure 2 and uses MITRA to automatically synthesize the desired data migration program. We now give a brief overview of how MITRA generates the target program.

Internally, MITRA represents the input XML file as a *hierarchical data tree*, shown in Figure 4a. Here, each node corresponds to an element in the XML document, and an edge from n to n' indicates that element n' is inside n . The bold text in each node from Figure 4a corresponds to the data stored in that element.

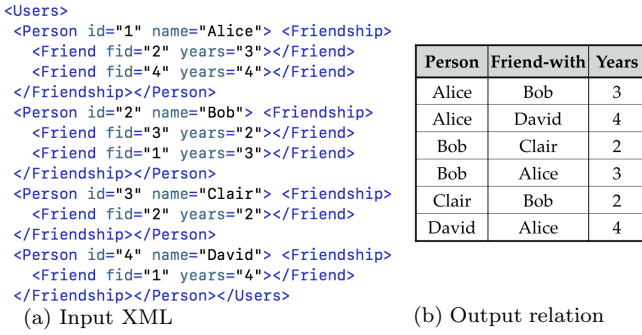


Figure 2: Input-output example

As mentioned in Section 1, MITRA starts by learning all possible *column extraction programs* that can be used to obtain column i in the output table from the input tree. Figure 3 shows the extraction programs for each column. Specifically, MITRA learns a single column extraction program π_{11} (resp. π_{21}) for the column `Person` (resp. column `Friend-with`). For instance, the program π_{11} first retrieves all children with tag `Person` of the root node, and, for each of these children, it returns the child with tag `name`. Since there are several ways to obtain the data in the `years` column, MITRA learns four different column extractors (namely, $\pi_{31}, \dots, \pi_{34}$) for `years`.

Next, MITRA conceptually generates intermediate tables by applying each column extractor to the input tree and taking their cross product.² Since we have four different column extraction programs for the `years` attribute, MITRA considers four different intermediate tables, one of which is shown in Figure 4b. In particular, this table is obtained using the *table extraction* program ψ presented in Figure 3. Observe that entries in the intermediate tables generated by MITRA refer to nodes from the input tree.

In the second phase of the synthesis algorithm, MITRA filters out spurious tuples in the intermediate table by learning a suitable *predicate*. For instance, the intermediate table from Figure 4b contains several tuples that do not appear in the output table. As an example, consider the tuple (n_7, n_7, n_{25}) . If we extract the data stored at these nodes, we would obtain the tuple $(\text{Alice}, \text{Alice}, 3)$, which is not part of the desired output table. To filter out such spurious tuples from the intermediate table, MITRA learns a predicate ϕ such that ϕ evaluates to *true* for every tuple in the target table, and evaluates to *false* for the spurious tuples. For our running example and the intermediate table from Figure 4b, MITRA learns the predicate $\phi_1 \wedge \phi_2$, where ϕ_1 and ϕ_2 are shown in Figure 3. Here, ϕ_1 ensures that a tuple (a, b, c) only exists if a and c are associated with the same person. Similarly, ϕ_2 guarantees that b refers to the person who has been friends with a for c years. For instance, since ϕ_2 evaluates to false for the first tuple in the intermediate table from Figure 4b, this spurious tuple will be filtered out by the learnt predicate.

While all programs synthesized by MITRA are guaranteed to satisfy the input-output examples, not all of these programs may correspond to the user’s intent. In particular, since examples are typically under-constrained, there may

²Since this strategy may be inefficient, our implementation performs an optimization to eliminate the generation of intermediate tables in most cases. However, decomposing the problem in this manner greatly facilitates the synthesis task.

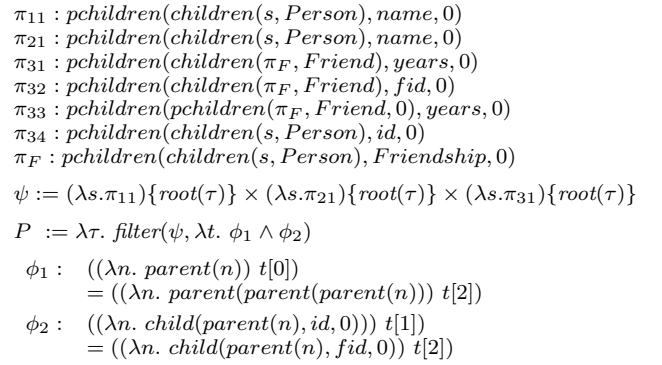


Figure 3: Synthesized program for the motivating example.

be multiple programs that satisfy the provided examples. For instance, in our running example, MITRA finds four different programs that are consistent with the examples. Our synthesis algorithm uses a heuristic ranking function based on *Occam’s razor* principle to rank-order the synthesized programs. For instance, programs that involve complex predicates are assigned a lower score by the ranking algorithm. Since the program P shown in Figure 3 has the highest rank, MITRA chooses it as the desired solution and optimizes P by memoizing redundant computations and avoiding the generation of intermediate tables whenever possible. Finally, MITRA generates an executable XSLT program, which is available from `goo.gl/rC4HT4`. The user can now apply the generated code to the original (much larger) XML document and obtain the desired relational representation. For instance, the synthesized program can be used to migrate an XML document with more than 1 million elements to the desired relational table in 154 seconds.

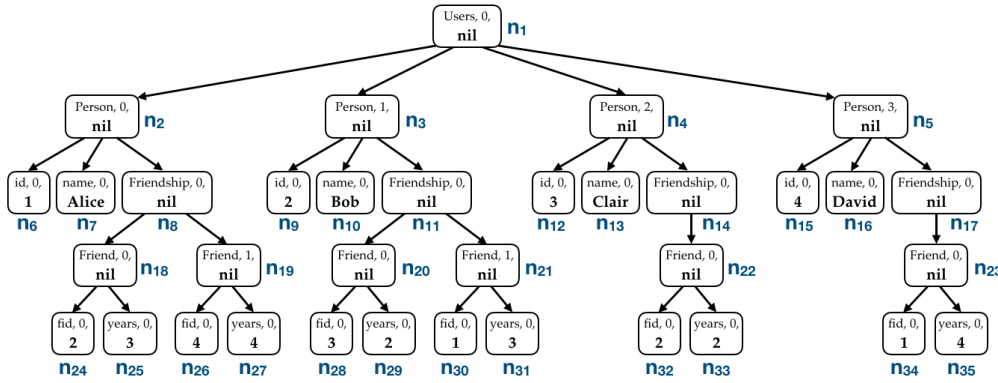
3. PRELIMINARIES

To allow for a uniform representation, we model tree-structured documents, such as XML, HTML and JSON, using so-called *hierarchical data trees*, which is a variant of the data structure used in [56].

DEFINITION 1. (Hierarchical Data Tree) *A hierarchical data tree (HDT) τ is a rooted tree represented as a tuple $\langle V, E \rangle$ where V is a set of nodes, and E is a set of directed edges. A node $v \in V$ is a triple $(tag, pos, data)$ where tag is a label associated with node v , pos indicates that v is the pos’th element with label tag under v ’s parent, and data corresponds to the data stored at node v .*

Given a node $n = (t, i, d)$ in a hierarchical data tree, we use the notation $n.tag$, $n.pos$, and $n.data$ to indicate t , i , and d respectively. We also use the notation $isLeaf(n)$ to denote that n is a leaf node of the HDT. In the rest of this paper, we assume that only leaf nodes contain data; hence, for any internal node (t, i, d) , we have $d = nil$. Finally, given an HDT τ , we write $\tau.root$ to denote the root node of τ .

XML documents as HDTs. We can represent XML files as hierarchical data trees where nodes correspond to XML elements. In particular, an edge from node v' to $v = (t, i, d)$ indicates that the XML element e represented by v is nested directly inside element e' represented by v' . Furthermore since v has tag t and position i , this means e is the i ’th child with tag t of e' . We also model XML attributes and text content as nested elements. This design choice allows



(a) Hierarchical data tree representation of the input XML

| $\pi_{11}(\{n_1\})$ | $\pi_{21}(\{n_1\})$ | $\pi_{31}(\{n_1\})$ |
|---------------------|---------------------|---------------------|
| n_7 | n_7 | n_{25} |
| n_7 | n_7 | n_{27} |
| n_7 | n_7 | n_{29} |
| n_7 | n_7 | n_{31} |
| n_7 | n_7 | n_{33} |
| n_7 | n_7 | n_{35} |
| n_7 | n_{10} | n_{25} |
| n_7 | n_{10} | n_{27} |
| ... | ... | ... |
| n_{16} | n_{16} | n_{31} |
| n_{16} | n_{16} | n_{33} |
| n_{16} | n_{16} | n_{35} |

(b) Intermediate table

Figure 4: Column extraction in the motivating example

```

{"Users": {"Person": [
  {"id": 1, "name": "Alice", "Friendship":
    {"Friend": [{"fid": 2, "years": 3}, {"fid": 4, "years": 4}]},
  {"id": 2, "name": "Bob", "Friendship":
    {"Friend": [{"fid": 3, "years": 2}, {"fid": 1, "years": 3}]},
  {"id": 3, "name": "Clair", "Friendship":
    {"Friend": [{"fid": 2, "years": 2}]},
  {"id": 4, "name": "David", "Friendship":
    {"Friend": [{"fid": 1, "years": 4}]}}]}

```

Figure 5: Example of a JSON file

our model to represent elements that contain a combination of nested elements, attributes, and text content.

EXAMPLE 1. Figure 4a shows the HDT representation of the XML file from Figure 2a. Observe how attributes are represented as nested elements in the HDT representation.

JSON documents as HDTs. JSON documents store data as a set of nested key-value pairs. We can model JSON files as HDTs in the following way: Each node $n = (t, i, d)$ in the HDT corresponds to a key-value pair e in the JSON file, where t represents the key and d represents the value. Since values in JSON files can be arrays, the position i corresponds to the i 'th entry in the array. For instance, if the JSON file maps key k to the array $[18, 45, 32]$, the HDT contains three nodes of the form $(k, 0, 18)$, $(k, 1, 45)$, $(k, 2, 32)$. An edge from n' to n indicates that the key-value pair e represented by n is nested inside the key-value pair e' represented by n' .

EXAMPLE 2. Figure 5 shows the JSON document corresponding to the HDT representation in Figure 4a. Observe that JSON objects and arrays are represented as internal nodes with data nil. For a given node n , we have $n.pos = 0$ unless the parent of n corresponds to an array.

4. DOMAIN-SPECIFIC LANGUAGE

In this section, we present our domain-specific language (DSL) for implementing transformations from hierarchical data trees to relational tables. As standard, we represent relational tables as a bag of tuples, where each tuple is represented using a list. Given a relational table \mathcal{R} , we use the notation $column(\mathcal{R}, i)$ to denote the i 'th column of \mathcal{R} , and we use the terms “relation” and “table” interchangeably.

Figure 6 shows the syntax of our DSL, and Figure 7 gives its semantics. Before we explain the constructs in this DSL, an important point is that our language is designed to achieve a good trade-off between expressiveness and efficiency of synthesis. That is, while our DSL can express a large class

Program $P := \lambda\tau.filter(\psi, \lambda t.\phi)$
Table Extractor $\psi := (\lambda s.\pi) \{root(\tau)\} \mid \psi_1 \times \psi_2$
Column Extractor $\pi := s \mid children(\pi, tag)$
 $\mid pchildren(\pi, tag, pos)$
 $\mid descendants(\pi, tag)$
Predicate $\phi := ((\lambda n.\varphi) t[i]) \leq c$
 $\mid ((\lambda n.\varphi_1) t[i]) \leq ((\lambda n.\varphi_2) t[j])$
 $\mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg\phi$
Node Extractor $\varphi := n \mid parent(\varphi) \mid child(\varphi, tag, pos)$

Figure 6: Syntax of our DSL. Here, τ denotes the input tree, t is bound to a tuple of nodes in τ , and n denotes a node in τ . Furthermore, i and j are integers, and c is a constant value (string, integer, etc). $t[i]$ gives the i -th element in tuple t .

of tree-to-table transformations that arise in practice, it is designed to make automated synthesis practical.

The high-level structure of the DSL follows our synthesis methodology of decomposing the problem into two separate column and row extraction operations. In particular, a program P is of the form $\lambda\tau.filter(\psi, \lambda t.\phi)$, where τ is the input HDT and ψ is a *table extractor* that extracts a relation \mathcal{R}' from τ . As mentioned in Section 1, the extracted table \mathcal{R}' overapproximates the target table \mathcal{R} . Therefore, the top-level *filter* construct uses a predicate ϕ to filter out tuples in \mathcal{R}' that do not appear in \mathcal{R} .

A *table extractor* ψ constructs a table by taking the cartesian product of a number of columns, where an entry in each column is a “pointer” to a node in the input HDT. Each column is obtained by applying a *column extractor* π to the root node of the input tree. A column extractor π takes as input a set of nodes and an HDT, and returns a set of HDT nodes. Column extractors are defined recursively and can be nested inside each other: The base case simply returns the input set of nodes, and the recursive case extracts other nodes from each input node using the *children*, *pchildren*, and *descendants* constructs. Specifically, the *children* construct extracts all children with a given tag, whereas *pchildren* yields all children with a given tag and specified position. In contrast, the *descendants* construct returns all descendants with the given tag. The formal (denotational) semantics of each construct is given in Figure 7.

Let us now turn our attention to predicates ϕ that can be used in the top-level *filter* construct. Atomic predicates without boolean connectives are obtained by comparing the

$$\begin{aligned}
\llbracket \text{filter}(\psi, \lambda t. \phi) \rrbracket_{\mathbb{T}} &= \{ (n_1.\text{data}, \dots, n_k.\text{data}) \mid t \in \llbracket \psi \rrbracket_{\mathbb{T}}, t = (n_1, \dots, n_k), \llbracket \phi \rrbracket_{t, \mathbb{T}} = \text{True} \} \\
\llbracket \psi_1 \times \psi_2 \rrbracket_{\mathbb{T}} &= \{ (n_1, \dots, n_k, n'_1, \dots, n'_l) \mid (n_1, \dots, n_k) \in \llbracket \psi_1 \rrbracket_{\mathbb{T}}, (n'_1, \dots, n'_l) \in \llbracket \psi_2 \rrbracket_{\mathbb{T}} \} \\
\llbracket (\lambda s. \pi) \{\text{root}(\tau)\} \rrbracket_{\mathbb{T}} &= \{ n \mid n \in \llbracket \pi \rrbracket_{s, \mathbb{T}}, s = \{\text{T.root}\} \} \\
\llbracket x \rrbracket_{s, \mathbb{T}} &= s \quad \text{where } s \text{ is a set of nodes in } \mathbb{T} \\
\llbracket \text{children}(\pi, \text{tag}) \rrbracket_{s, \mathbb{T}} &= \{ n' \mid n \in \llbracket \pi \rrbracket_{s, \mathbb{T}}, (n, n') \in E, n'.\text{tag} = \text{tag} \} \text{ where } \mathbb{T} = (V, E) \\
\llbracket \text{pchildren}(\pi, \text{tag}, \text{pos}) \rrbracket_{s, \mathbb{T}} &= \{ n' \mid n \in \llbracket \pi \rrbracket_{s, \mathbb{T}}, (n, n') \in E, n'.\text{tag} = \text{tag}, n'.\text{pos} = \text{pos} \} \text{ where } \mathbb{T} = (V, E) \\
\llbracket \text{descendants}(\pi, \text{tag}) \rrbracket_{s, \mathbb{T}} &= \{ n_z \mid n_1 \in \llbracket \pi \rrbracket_{s, \mathbb{T}}, \exists \{n_2, \dots, n_{z-1}\} \subset V \text{ s.t. } \forall 1 \leq x < z. (n_x, n_{x+1}) \in E, n_z.\text{tag} = \text{tag} \} \text{ where } \mathbb{T} = (V, E) \\
\llbracket ((\lambda n. \varphi) t[i]) \leq c \rrbracket_{t, \mathbb{T}} &= n'.\text{data} \leq c \quad \text{where } t = (n_1, \dots, n_l) \text{ and } n' = \llbracket \varphi \rrbracket_{n_i, \mathbb{T}} \\
\llbracket ((\lambda n. \varphi_1) t[i]) \leq ((\lambda n. \varphi_2) t[j]) \rrbracket_{t, \mathbb{T}} &= \begin{cases} n'_1.\text{data} \leq n'_2.\text{data} & \text{if } n'_1 \text{ and } n'_2 \text{ are both leaf nodes of } \mathbb{T} \\ n'_1 = n'_2 & \text{if } \leq \text{ is "="}, \text{ and neither } n'_1 \text{ nor } n'_2 \text{ is a leaf node of } \mathbb{T} \\ \text{False} & \text{otherwise} \end{cases} \\
&\quad \text{where } t = (n_1, \dots, n_l) \text{ and } n'_1 = \llbracket \varphi_1 \rrbracket_{n_i, \mathbb{T}} \text{ and } n'_2 = \llbracket \varphi_2 \rrbracket_{n_j, \mathbb{T}} \\
\llbracket \phi_1 \wedge \phi_2 \rrbracket_{t, \mathbb{T}} &= \llbracket \phi_1 \rrbracket_{t, \mathbb{T}} \wedge \llbracket \phi_2 \rrbracket_{t, \mathbb{T}} \\
\llbracket \phi_1 \vee \phi_2 \rrbracket_{t, \mathbb{T}} &= \llbracket \phi_1 \rrbracket_{t, \mathbb{T}} \vee \llbracket \phi_2 \rrbracket_{t, \mathbb{T}} \\
\llbracket \neg \phi \rrbracket_{t, \mathbb{T}} &= \neg \llbracket \phi \rrbracket_{t, \mathbb{T}} \\
\llbracket n \rrbracket_{n, \mathbb{T}} &= n \\
\llbracket \text{parent}(\varphi) \rrbracket_{n, \mathbb{T}} &= \begin{cases} n' & \text{if } (n', \llbracket \varphi \rrbracket_{n, \mathbb{T}}) \in E \\ \perp & \text{otherwise} \end{cases} \quad \text{where } \mathbb{T} = (V, E) \\
\llbracket \text{child}(\varphi, \text{tag}, \text{pos}) \rrbracket_{n, \mathbb{T}} &= \begin{cases} n' & \text{if } (\llbracket \varphi \rrbracket_{n, \mathbb{T}}, n') \in E \text{ and } n'.\text{tag} = \text{tag} \text{ and } n'.\text{pos} = \text{pos} \\ \perp & \text{otherwise} \end{cases} \quad \text{where } \mathbb{T} = (V, E)
\end{aligned}$$

Figure 7: Semantics of our DSL.

```

<Dependency> <Object id="8"> A
  <Object id="13"> B <Object id="18"> D
    <Object id="23"> F </Object>
    <Object id="28"> G </Object> </Object> </Object>
  <Object id="20"> C <Object id="24"> E </Object> </Object>
</Object> </Dependency>

```

(a) Input XML

| Parent | Child |
|--------|-------|
| A | B |
| A | C |
| B | D |
| D | F |
| D | G |

(b) Output relation

```

P := λτ. filter(ψ, λt. φ1 ∧ φ2)
ψ := (λs. π1) {root(τ)} × (λs. π2) {root(τ)}
π1 = π2 = pchildren(descendants(s, Object), text, 0)
φ1 := ((λn. child(parent(n), id, 0)) t[0]) < 20
φ2 := ((λn. parent(n)) t[0]) = ((λn. parent(parent(n)) t[1])

```

(c) Synthesized program

Figure 8: Input-output example and the synthesized program for Example 3

data stored in an HDT node with a constant c or the data stored in another tree node. In particular, predicates make use of *node extractors* φ that take a tree node as input and return another tree node. Similar to column extractors, node extractors are recursively defined and can be nested within each other. Observe that node extractors allow accessing both parent and children nodes; hence, they can be used to extract any arbitrary node within the HDT from a given input node. (Figure 7 gives the formal semantics).

Going back to the definition of predicates, ϕ takes a tuple t and evaluates to a boolean value indicating whether t should be kept in the output table. The simplest predicate $((\lambda n. \varphi) t[i]) \leq c$ first extracts the i 'th entry n of t and then uses the node extractor φ to obtain another tree node n' from n . This predicate evaluates to true iff $n'.\text{data} \leq c$ is true. The semantics of $((\lambda n. \varphi_1) t[i]) \leq ((\lambda n. \varphi_2) t[j])$ is similar, except that it compares values stored at two tree nodes n, n' . In particular, if n, n' are both leaf nodes, then we check whether the relation $n.\text{data} \leq n'.\text{data}$ is satisfied. If they are internal nodes and the operator \leq is equality, then we check whether n, n' refer to the same node. Otherwise, the predicate evaluates to false. More complex predicates are obtained using the standard boolean connectives \wedge, \vee, \neg .

EXAMPLE 3. Consider the data transformation task illustrated in Figure 8, in which we want to map the text value of each object element with id less than 20 in the XML file to the text value of its immediate nested object elements. Figure 8c shows the DSL program for this transformation. Here, column extractors π_1, π_2 use the descendants and pchildren constructs to extract all children nodes with tag text and pos 0 of any object node reachable from the root. Predicate ϕ_1

filters out all tuples where the first element in the tuple has an id sibling with value greater than or equal to 20. The second predicate ϕ_2 ensures that the second element in the tuple is directly nested inside the first one.

5. SYNTHESIS ALGORITHM

In this section, we present our synthesis algorithm for converting an HDT into a relational table from input-output examples. Our technique can be used to transform an XML or JSON document into a relational database by running the algorithm once for each table in the target database.

The top-level structure of our synthesis algorithm is given in Algorithm 1. The algorithm takes a set of input-output examples of the form $\{\mathbb{T}_1 \rightarrow \mathcal{R}_1, \dots, \mathbb{T}_m \rightarrow \mathcal{R}_m\}$, where each \mathbb{T}_i represents an HDT (input example) and \mathcal{R}_i represents its desired relational table representation (output example). Since the schema of the target table is typically fixed in practice, we assume that all output tables (i.e., \mathcal{R}_i 's) have the same number of columns. Given these input-output examples, the procedure `LEARNTRANSFORMATION` returns a program P^* in our DSL that is consistent with all input-output examples. Furthermore, since P^* is the simplest DSL program that satisfies the specification, it is expected to be a general program that is not over-fitted to the examples.

As mentioned earlier, our methodology decomposes the synthesis task into two phases for extracting columns and rows. In the first phase, we synthesize a *column extraction* program that yields an overapproximation of each column in the output table \mathcal{R} . The cartesian product of columns extracted by these column extraction programs further yields a table \mathcal{R}' that overapproximates the target table \mathcal{R} . In the

Algorithm 1 Top-level synthesis algorithm

```
1: procedure LEARNTRANSFORMATION( $\mathcal{E}$ )
2:   Input: Examples  $\mathcal{E} = \{\tau_1 \rightarrow \mathcal{R}_1, \dots, \tau_m \rightarrow \mathcal{R}_m\}$ .
3:   Output: Simplest DSL program  $P^*$ .
4:   for all  $1 \leq j \leq k$  do
5:      $\Pi_j := \text{LEARNCOLEXTRACTORS}(\mathcal{E}, j)$ ;
6:    $\Psi := \Pi_1 \times \dots \times \Pi_n$ ;
7:    $P^* := \perp$ ;
8:   for all  $\psi \in \Psi$  do
9:      $\phi := \text{LEARNPREDICATE}(\mathcal{E}, \psi)$ ;
10:    if  $\phi \neq \perp$  then
11:       $P := \lambda\tau. \text{filter}(\psi, \lambda t. \phi)$ ;
12:      if  $\theta(P) < \theta(P^*)$  then  $P^* := P$ 
13:   return  $P^*$ ;
```

second phase, we learn a predicate ϕ that allows us to filter out exactly the “spurious” tuples in \mathcal{R}' that do not appear in the output table \mathcal{R} .

Let us now consider the LEARNTRANSFORMATION procedure from Algorithm 1 in more detail. Given the examples, we first invoke a procedure called LEARNCOLEXTRACTORS that learns a set Π_j of column extraction expressions such that applying any $\pi \in \Pi_j$ on each τ_i yields a column that overapproximates the j 'th column in table \mathcal{R}_i (lines 4-5). Observe that our algorithm learns a set of column extractors (instead of just one) because some of them might not lead to a desired *filter* program. Our procedure for learning column extractors is based on deterministic finite automata and will be described in more detail in Section 5.1.

Once we have the individual column extraction programs for each column, we then obtain the set of all possible table extraction programs by taking the cartesian product of each column extraction program (line 6). Hence, applying each table extraction program $\psi \in \Psi$ to the input tree τ_i yields a table \mathcal{R}'_i that overapproximates the output table \mathcal{R}_i .

The next step of the synthesis algorithm (lines 7-12) learns the predicate used in the top-level *filter* construct. For each table extraction program $\psi \in \Psi$, we try to learn a predicate ϕ such that $\lambda\tau. \text{filter}(\psi, \lambda t. \phi)$ is consistent with the examples. Specifically, the procedure LEARNPREDICATE yields a predicate that allows us to filter out spurious tuples in $\llbracket \psi \rrbracket_{\tau_i}$. If there is no such predicate, LEARNPREDICATE returns \perp . Our predicate learning algorithm uses integer linear programming to infer a *simplest* formula with the minimum number of atomic predicates. We describe the LEARNPREDICATE algorithm in more detail in Section 5.2.

Since there may be multiple DSL programs that satisfy the provided examples, our method uses the Occam’s razor principle to choose between different solutions. In particular, Algorithm 1 uses a heuristic cost function θ to determine which solution is simpler (line 12), so it is guaranteed to return a program that minimizes the value of θ . Intuitively, θ assigns a higher cost to programs that use complex predicates or column extractors. We discuss the design of the heuristic cost function θ in Section 6.

5.1 Learning Column Extraction Programs

Our technique for learning column extraction programs is based on deterministic finite automata (DFA). Given a tree (input example) and a single column (output example), our method constructs a DFA whose language accepts exactly the set of column extraction programs that are consistent

Algorithm 2 Algorithm for learning column extractors

```
1: procedure LEARNCOLEXTRACTORS( $\mathcal{E}, i$ )
2:   Input: Examples  $\mathcal{E}$  and column number  $i$ .
3:   Output: Set  $\Pi$  of column extractors.
4:    $\mathcal{E}' := \{(\tau, \kappa) \mid (\tau, \mathcal{R}) \in \mathcal{E} \wedge \kappa = \text{column}(\mathcal{R}, i)\}$ 
5:    $\mathcal{A} := \text{CONSTRUCTDFA}(e_0)$  where  $e_0 \in \mathcal{E}'$ 
6:   for all  $e \in \mathcal{E}' \setminus \{e_0\}$  do
7:      $\mathcal{A}' := \text{CONSTRUCTDFA}(e)$ 
8:      $\mathcal{A} := \text{INTERSECT}(\mathcal{A}, \mathcal{A}')$ 
9:   return  $\text{LANGUAGE}(\mathcal{A})$ 
```

$$\frac{s = \{\tau.\text{root}\}}{q_0 = q_s \in Q} \quad (1)$$

$$\frac{q_s \in Q \quad \text{tag is a tag in } \tau \quad \llbracket \text{children}(s, \text{tag}) \rrbracket_{s, \tau} = s'}{q_{s'} \in Q, \delta(q_s, \text{children}_{\text{tag}}) = q_{s'}} \quad (2)$$

$$\frac{q_s \in Q \quad \text{tag is a tag in } \tau \quad \text{pos is a pos in } \tau \quad \llbracket \text{pchildren}(s, \text{tag}, \text{pos}) \rrbracket_{s, \tau} = s'}{q_{s'} \in Q, \delta(q_s, \text{pchildren}_{\text{tag}, \text{pos}}) = q_{s'}} \quad (3)$$

$$\frac{q_s \in Q \quad \text{tag is a tag in } \tau \quad \llbracket \text{descendants}(s, \text{tag}) \rrbracket_{s, \tau} = s'}{q_{s'} \in Q, \delta(q_s, \text{descendants}_{\text{tag}}) = q_{s'}} \quad (4)$$

$$\frac{s \supseteq \text{column}(\mathcal{R}, i)}{q_s \in F} \quad (5)$$

Figure 9: FTA construction rules. τ is the input tree, \mathcal{R} is the output table, and i is the column to be extracted.

with the input-output example. If we have multiple input-output examples, we construct a separate DFA for each example and then take their intersection. The language of the resulting DFA includes column extraction programs that are consistent with *all* examples.

Let us now look at the LEARNCOLEXTRACTORS procedure shown in Algorithm 2 in more detail. It takes the set of input-output examples \mathcal{E} as well as a column number i for which we would like to learn the extractor. The algorithm returns a set of column extraction programs Π such that for every $\pi \in \Pi$ and every input-output example (τ, \mathcal{R}) , we have $\llbracket \pi \rrbracket_{\{\tau.\text{root}\}, \tau} \supseteq \text{column}(\mathcal{R}, i)$.

Algorithm 2 starts by constructing a set of examples \mathcal{E}' mapping each input tree to the i 'th column of the output table (line 4). Then, for each example $e \in \mathcal{E}'$, we construct a DFA that represents all column extraction programs consistent with e (lines 5-8). The set of programs consistent with all examples \mathcal{E}' corresponds to the language of the intersection automaton \mathcal{A} from line 9. In particular, the INTERSECT procedure used in line 8 is the standard intersection procedure for DFAs [28]. Concretely, the intersection of two DFAs \mathcal{A}_1 and \mathcal{A}_2 only accepts programs that are accepted by both \mathcal{A}_1 and \mathcal{A}_2 and is constructed by taking the cartesian product of \mathcal{A}_1 and \mathcal{A}_2 and defining the accepting states to be all states of the form (q_1, q_2) where q_1 and q_2 are accepting states of \mathcal{A}_1 and \mathcal{A}_2 respectively.

The key part of the LEARNCOLEXTRACTORS procedure is the CONSTRUCTDFA method, which constructs a deterministic finite automaton $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ from an input-output example using the rules shown in Figure 9. Here,

the states Q of the automaton correspond to sets of nodes in the input HDT. We use the notation q_s to denote the state representing the set of nodes s . The alphabet Σ corresponds to the names of column extraction functions in the DSL. Specifically, we have:

$$\begin{aligned} \Sigma &= \{ \text{children}_{\text{tag}} \mid \text{tag is a tag in } \mathbb{T} \} \\ &\cup \{ \text{pchildren}_{\text{tag, pos}} \mid \text{tag (pos) is a tag (pos) in } \mathbb{T} \} \\ &\cup \{ \text{descendants}_{\text{tag}} \mid \text{tag is a tag in } \mathbb{T} \} \end{aligned}$$

In other words, each symbol in the alphabet corresponds to a DSL operator (instantiated with labels and positions from the input HDT). Transitions in the DFA are constructed using the semantics of DSL operators: Intuitively, given a DSL construct $f \in \{ \text{children}, \text{pchildren}, \text{descendants} \}$ and a state q_s , the DFA contains a transition $q_s \xrightarrow{f} q'_s$ if applying f to s produces s' . The initial state of the DFA is $\{\mathbb{T}.root\}$, and we have $q_s \in F$ iff s overapproximates the i 'th column in table \mathcal{R} .

Let us now look at the construction rules shown in Figure 9 in more detail. Rules (1)-(4) process the column extractor constructs in our DSL and construct states and/or transitions. The first rule adds $q_{\{\mathbb{T}.root\}}$ as an initial state because the root node of the HDT is directly reachable. The second rule adds a state $q_{s'}$ and a transition $\delta(q_s, \text{children}_{\text{tag}}) = q_{s'}$ if $\text{children}(s, \text{tag})$ evaluates to s' . Rules (3) and (4) are similar to rule (2) and process the remaining column extraction functions pchildren and descendants . For example, we have $\delta(q_s, \text{pchildren}_{\text{tag, pos}}) = q_{s'}$ if $\text{pchildren}(s, \text{tag, pos})$ evaluates to s' . The last rule in Figure 9 identifies the final states. In particular, a state q_s is a final state if s is a superset of the target column (i.e., output example).

THEOREM 1. *Let \mathcal{A} be the DFA constructed by Algorithm 2 for a set of input-output examples \mathcal{E} and a column number i . Then, \mathcal{A} accepts a column extraction program π in our DSL iff $\forall (\mathbb{T}, \mathcal{R}) \in \mathcal{E}. \llbracket \pi \rrbracket_{\{\mathbb{T}.root\}, \mathbb{T}} \supseteq \text{column}(\mathcal{R}, i)$.*³

EXAMPLE 4. *Suppose the DFA constructed using rules from Figure 9 accepts the word ab where $a = \text{descendants}_{\text{object}}$ and $b = \text{pchildren}_{\text{text}, 0}$. This word corresponds to the following column extractor program:*

$$\pi = \text{pchildren}(\text{descendants}(s, \text{object}), \text{text}, 0)$$

If the input example is \mathbb{T} and the output example is $\text{column}(\mathcal{R}, i)$, then we have $((\lambda s.\pi) \{\mathbb{T}.root\}) \supseteq \text{column}(\mathcal{R}, i)$.

5.2 Learning Predicates

We now turn our attention to the predicate learning algorithm `LEARNPREDICATE` shown in Algorithm 3. This procedure takes the input-output examples \mathcal{E} and a candidate table extractor ψ and returns a predicate ϕ such that for every $(\mathbb{T}, \mathcal{R}) \in \mathcal{E}$, the program $\text{filter}(\psi, \lambda t.\phi)$ yields the desired output table \mathcal{R} on input \mathbb{T} .

The algorithm starts by constructing a (finite) universe Φ of all possible atomic predicates that can be used in formula ϕ (line 4). These predicates are constructed for a set of input-output examples \mathcal{E} and a candidate table extractor ψ using rules from Figure 10. While these rules are not very important for understanding the key ideas of our technique, let us consider rule (4) from Figure 10 as an example. According to this rule, we generate an atomic predicate

³Proofs of all theorems are available in the extended version of this paper in [57].

Algorithm 3 Algorithm for learning predicates

```

1: procedure LEARNPREDICATE( $\mathcal{E}, \psi$ )
2:   Input: Examples  $\mathcal{E}$ , a candidate table extractor  $\psi$ .
3:   Output: Desired predicate  $\phi$ .
4:    $\Phi := \text{CONSTRUCTPREDUNIVERSE}(\mathcal{E}, \psi)$ 
5:    $\mathcal{E}^+ := \emptyset; \mathcal{E}^- := \emptyset$ 
6:   for all  $(\mathbb{T}, \mathcal{R}) \in \mathcal{E}$  do
7:     for all  $t \in \llbracket \psi \rrbracket_{\mathbb{T}}$  do
8:       if  $t \in \mathcal{R}$  then
9:          $\mathcal{E}^+ := \mathcal{E}^+ \cup \{t\}$ 
10:      else  $\mathcal{E}^- := \mathcal{E}^- \cup \{t\}$ 
11:    $\Phi^* := \text{FINDMINCOVER}(\Phi, \mathcal{E}^+, \mathcal{E}^-)$ 
12:   Find  $\phi$  such that:
13:     (1)  $\phi$  is boolean combination of preds in  $\Phi^*$ 
14:     (2)  $\phi(t) = \begin{cases} 1 & \text{if } t \in \mathcal{E}^+ \\ 0 & \text{if } t \in \mathcal{E}^- \end{cases}$ 
15:   return  $\phi$ 

```

Algorithm 4 Algorithm to find minimum predicate set

```

1: procedure FINDMINCOVER( $\Phi, \mathcal{E}^+, \mathcal{E}^-$ )
2:   Input: Universe of predicates  $\Phi$ 
3:   Input: Positive examples  $\mathcal{E}^+$ , negative examples  $\mathcal{E}^-$ 
4:   Output: Set of predicates  $\Phi^*$  where  $\Phi^* \subseteq \Phi$ 
5:   for all  $(\phi_k, e_i, e_j) \in \Phi \times \mathcal{E}^+ \times \mathcal{E}^-$  do
6:      $a_{ijk} = \begin{cases} 1 & \text{if } \phi_k(e_i) \neq \phi_k(e_j) \\ 0 & \text{otherwise} \end{cases}$ 
7:      $\text{minimize } \sum_{k=1}^{|\Phi|} x_k$ 
8:     subject to:
9:      $\forall (e_i, e_j) \in \mathcal{E}^+ \times \mathcal{E}^- . \sum_{k=1}^{|\Phi|} a_{ijk} \cdot x_k \geq 1$ 
10:     $\wedge \forall k \in [1, |\Phi|]. x_k \in \{0, 1\}$ 
11:   return  $\{\phi_i \mid \phi_i \in \Phi \wedge x_i = 1\}$ 

```

$((\lambda n.\varphi) t[i]) \leq c$ if i is a valid column number in the range $[1, k]$, c is a constant in one of the input tables, and φ is a “valid” node extractor for column i . Rules (1)-(3) in Figure 10 define what it means for a node extractor φ to be valid for column i , denoted as $\varphi \in \chi_i$. In particular, we say φ is a valid node extractor for column i if applying φ does not “throw an exception” (i.e., yield \perp) for any of the entries in the i 'th column of the generated intermediate tables.⁴

The next step of `LEARNPREDICATE` constructs a set of *positive* and *negative* examples to be used in `LEARNPREDICATE` (lines 5–10). In this context, positive examples \mathcal{E}^+ refer to tuples that should be present in the desired output table, whereas negative examples \mathcal{E}^- correspond to tuples that should be filtered out. The goal of the predicate learner is to find a formula ϕ over atomic predicates in Φ such that ϕ evaluates to true for all positive examples and to false for all negative examples. In other words, formula ϕ serves as a classifier between \mathcal{E}^+ and \mathcal{E}^- .

To learn a suitable classifier, our algorithm first learns a *minimum* set of atomic predicates that are necessary for distinguishing the \mathcal{E}^+ samples from the \mathcal{E}^- ones. Since our goal is to synthesize a general program that is not overfitted to the input-output examples, it is important that the synthesized predicate ϕ is as simple as possible. We

⁴Recall that these intermediate tables are obtained by applying ψ to the input HDTs in \mathcal{E} .

$$\frac{}{\pi_i, \mathcal{E} \vdash n \in \chi_i} \quad (1)$$

$$\frac{\pi_i, \mathcal{E} \vdash \varphi \in \chi_i}{\forall T \rightarrow \mathcal{R} \in \mathcal{E}. \forall n \in \pi_i(T). \llbracket \text{parent}(\varphi) \rrbracket_{n,T} \neq \perp} \quad (2)$$

$$\frac{\pi_i, \mathcal{E} \vdash \varphi \in \chi_i}{\forall T \rightarrow \mathcal{R} \in \mathcal{E}. \forall n \in \pi_i(T). \llbracket \text{child}(\varphi, \text{tag}, \text{pos}) \rrbracket_{n,T} \neq \perp} \quad (3)$$

$$\frac{\begin{array}{l} \psi = \pi_1 \times \dots \times \pi_k, \quad i \in [1, k] \\ \pi_i, \mathcal{E} \vdash \varphi \in \chi_i \\ \exists (T, \mathcal{R}) \in \mathcal{E}. c \in \text{data}(T) \end{array}}{\psi, \mathcal{E} \vdash ((\lambda n. \varphi) t[i]) \trianglelefteq c \in \Phi} \quad (4)$$

$$\frac{\begin{array}{l} \psi = \pi_1 \times \dots \times \pi_k, \quad i \in [1, k], \quad j \in [1, k] \\ \pi_i, \mathcal{E} \vdash \varphi_1 \in \chi_i \\ \pi_j, \mathcal{E} \vdash \varphi_2 \in \chi_j \end{array}}{\psi, \mathcal{E} \vdash ((\lambda n. \varphi_1) t[i]) \trianglelefteq ((\lambda n. \varphi_2) t[j]) \in \Phi} \quad (5)$$

Figure 10: Predicate universe construction rules. \mathcal{E} is the input-output examples and $\psi = \pi_1 \times \dots \times \pi_k$ is the candidate table extractor. Here χ_i indicates a set of node extractors that can be applied to the nodes extracted for column i .

formulate the problem of finding a simplest classifier as a combination of *integer linear programming (ILP)* and *logic minimization* [37, 42]. In particular, we use ILP to learn a minimum set of predicates Φ^* that must be used in the classifier, and then use circuit minimization techniques to find a DNF formula ϕ over Φ^* with the minimum number of boolean connectives.

Our method for finding the minimum set of atomic predicates is given in Algorithm 4. The FINDMINCOVER procedure takes as input the universe Φ of all predicates as well as positive and negative examples \mathcal{E}^+ , \mathcal{E}^- . It returns a subset Φ^* of Φ such that, for every pair of examples $(e_1, e_2) \in \mathcal{E}^+ \times \mathcal{E}^-$, there exists an atomic predicate $p \in \Phi^*$ such that p evaluates to different truth values for e_1 and e_2 (i.e., p differentiates e_1 and e_2).

We solve this optimization problem by reducing it to 0-1 ILP in the following way: First, we introduce an indicator variable x_k such that $x_k = 1$ if $p_k \in \Phi$ is chosen to be in Φ^* and $x_k = 0$ otherwise. Then, for each predicate p_k and every pair of examples $(e_i, e_j) \in \mathcal{E}^+ \times \mathcal{E}^-$, we introduce a (constant) variable a_{ijk} such that we assign a_{ijk} to 1 if predicate p_k distinguishes e_i and e_j and to 0 otherwise. Observe that the value of each a_{ijk} is known, whereas the assignment to each variable x_k is to be determined.

To find an optimal assignment to variables \vec{x} , we set up the 0-1 ILP problem shown in lines 7–10 of Algorithm 4. First, we require that for every pair of examples, there exists a predicate p_k that distinguishes them. This requirement is captured using the constraint in line 9: Since a_{ijk} is 1 iff p_k distinguishes e_i and e_j , the constraint at line 9 is satisfied only when we assign at least one of the x_k ’s differentiating e_i and e_j to 1. The objective function at line 7 minimizes the sum of the x_k ’s, thereby forcing us to choose the minimum number of predicates that are sufficient to distinguish every pair of positive and negative examples.

Going back to Algorithm 3, the return value Φ^* of FINDMINCOVER (line 11) corresponds to a minimum set of predicates that must be used in the classifier; however, we still need to find a *boolean combination* of predicates in Φ^* that

| | ϕ_1 | ϕ_2 | ϕ_3 | ϕ_4 | ϕ_5 | ϕ_6 | ϕ_7 |
|---------|----------|----------|----------|----------|----------|----------|----------|
| e_1^+ | true | true | false | false | true | true | false |
| e_2^+ | false | true | true | true | true | false | true |
| e_3^+ | false | true | true | true | false | false | false |
| e_1^- | false | false | true | true | false | false | false |
| e_2^- | false | true | true | true | false | false | true |
| e_3^- | true | false | true | false | false | false | true |

Figure 11: Initial truth table for the predicate universe Φ , positive examples \mathcal{E}^+ and negative examples \mathcal{E}^- .

| | v_{11} | v_{12} | v_{13} | v_{21} | v_{22} | v_{23} | v_{31} | v_{32} | v_{33} |
|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|
| ϕ_1 | 1 | 1 | 0 | 0 | 0 | 1 | 0 | 0 | 1 |
| ϕ_2 | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | 1 |
| ϕ_3 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 |
| ϕ_4 | 1 | 1 | 0 | 0 | 0 | 1 | 0 | 0 | 1 |
| ϕ_5 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 |
| ϕ_6 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 |
| ϕ_7 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 1 | 1 |

Figure 12: Values of a_{ijk} assigned in line 6 of Algorithm 4. Here v_{ij} corresponds to $(e_i, e_j) \in \mathcal{E}^+ \times \mathcal{E}^-$.

differentiates \mathcal{E}^+ samples from the \mathcal{E}^- ones. Furthermore, we would like to find the *smallest* such boolean combination for two reasons: (1) large formulas might hinder the generality of the synthesized program as well as its readability, and (2) large formulas would incur more overhead when being evaluated at runtime.

We cast the problem of finding a smallest classifier over predicates in Φ^* as a *logic minimization* problem [37, 42]. In particular, given a set of predicates Φ^* , our goal is to find a smallest DNF formula ϕ over predicates in Φ^* such that ϕ evaluates to *true* for any positive example and to *false* for any negative example. To solve this problem, we start by constructing a (partial) truth table, where the rows correspond to examples in $\mathcal{E}^+ \cup \mathcal{E}^-$, and the columns correspond to predicates in Φ^* . The entry in the i ’th row and j ’th column of the truth table is *true* if predicate $p_j \in \Phi^*$ evaluates to true for example e_i and false otherwise. The target boolean function ϕ should evaluate to true for any $e^+ \in \mathcal{E}^+$ and false for any $e^- \in \mathcal{E}^-$. Since we have a truth table describing the target boolean function, we can use standard techniques, such as the Quine-McCluskey method [37, 42], to find a smallest DNF formula representing classifier ϕ .

THEOREM 2. *Given examples \mathcal{E} and table extractor ψ such that $\forall (T, \mathcal{R}) \in \mathcal{E}. \mathcal{R} \subseteq \llbracket \psi \rrbracket_T$, Algorithm 3 returns a smallest DNF formula ϕ such that $\forall (T, \mathcal{R}) \in \mathcal{E}. \llbracket \text{filter}(\psi, \lambda t. \phi) \rrbracket_T = \mathcal{R}$ if such a formula exists in our DSL.*

EXAMPLE 5. *Consider predicate universe $\Phi = \{\phi_1, \dots, \phi_7\}$, a set of positive examples $\mathcal{E}^+ = \{e_1^+, e_2^+, e_3^+\}$, and a set of negative examples $\mathcal{E}^- = \{e_1^-, e_2^-, e_3^-\}$ with the truth table given in Figure 11. Here, the entry at row e_i and column ϕ_j of Figure 11 is true iff tuple e_i satisfies predicate ϕ_j . The goal of the predicate learner is to find a formula $\phi?$ with the minimum number of atomic predicates $\phi_i \in \Phi$ such that it evaluates to true for all positive examples and to false for all negative examples. In order to do so, the FINDMINCOVER procedure first finds the minimum required subset of atomic predicates Φ^* as described in Algorithm 4. Figure 12 shows the values of a_{ijk} assigned in line 6 of Algorithm 4. In particular, the entry at row ϕ_i and column v_{jk} of Figure 12 is true iff a_{ijk} is true. The highlighted rows in the Figure 12*

| | ϕ_2 | ϕ_5 | ϕ_7 | $\phi_?$ |
|---------|----------|----------|----------|----------|
| e_1^+ | true | true | false | true |
| e_2^+ | true | true | true | true |
| e_3^+ | true | false | false | true |
| e_1^- | false | false | false | false |
| e_2^- | true | false | true | false |
| e_3^- | false | false | true | false |

Figure 13: Truth table constructed in lines 14 – 16 of Algorithm 3 for Example 5.

indicate the predicates which are selected to be in Φ^* using integer linear programming. After finding $\Phi^* = \{\phi_2, \phi_5, \phi_7\}$, lines 12–14 Algorithm 3 generate the (partial) truth table as shown in Figure 13. The smallest DNF formula that is consistent with this truth table is $\phi_5 \vee (\phi_2 \wedge \neg\phi_7)$, so our algorithm returns this formula as a classifier.

The following two theorems state the soundness and completeness of our algorithm:

THEOREM 3. (Soundness) *Given examples $\mathcal{E} = \{T_1 \rightarrow \mathcal{R}_1, \dots, T_m \rightarrow \mathcal{R}_m\}$, suppose that $\text{LEARNTRANSFORMATION}(\mathcal{E})$ yields P^* . Then, $\forall (T_i, \mathcal{R}_i) \in \mathcal{E}$, we have $\llbracket P^* \rrbracket_{T_i} = \mathcal{R}_i$.*

THEOREM 4. (Completeness) *Suppose there is a DSL program consistent with examples $\mathcal{E} = \{T_1 \rightarrow \mathcal{R}_1, \dots, T_m \rightarrow \mathcal{R}_m\}$. Then, $\text{LEARNTRANSFORMATION}(\mathcal{E})$ eventually returns a program P^* such that $\forall i \in [1, m] : \llbracket P^* \rrbracket_{T_i} = \mathcal{R}_i$.*

Finally, the following theorem states that our synthesis algorithm returns a simplest DSL program with respect to the cost metric θ :

THEOREM 5. (Simplicity) *Given examples \mathcal{E} , Algorithm 1 returns a DSL program P^* such that for any program P satisfying \mathcal{E} , we have $\theta(P^*) \leq \theta(P)$.*

Complexity. Our algorithm has worst-case exponential time complexity with respect to the size of input-output examples, as *integer linear programming* and *logic minimization* for a given truth table are both NP-hard problems. However, *in practice*, the complexity of our algorithm does not come close to the worst-case scenario. A more detailed explanation of the empirical complexity of the proposed algorithm can be found in the extended version [57].

6. IMPLEMENTATION

We have implemented our synthesis algorithm in a tool called MITRA, which consists of approximately 7,000 lines of Java code. As shown in Figure 14, MITRA includes a domain-agnostic synthesis core (referred to as MITRA-CORE) and a set of domain-specific plug-ins. Specifically, MITRA-CORE accepts input-output examples in the form of (HDT, table) pairs and outputs a program over the DSL shown in Figure 6. The goal of a MITRA plug-in is to (1) convert the input document to our internal HDT representation, and (2) translate the program synthesized by MITRA-CORE to a target DSL. We have currently implemented two domain-specific plug-ins, called MITRA-XML and MITRA-JSON, for XML and JSON documents respectively. Specifically, MITRA-XML outputs programs written in XSLT, and MITRA-JSON generates JavaScript programs. MITRA can be easily extended to handle other forms of hierarchical documents (e.g., HTML and HDF) by implementing suitable plug-ins.




Figure 14: Architecture of MITRA

Cost function. Recall from Section 5 that our synthesis algorithm uses a heuristic cost function θ to choose the simplest program among all possible solutions that are consistent with the provided input-output examples. The cost function that we use in our implementation ranks programs based on the complexity of their predicates and column extractors and returns a program with the lowest cost. Given two programs P_1, P_2 , our cost function assigns P_1 (resp. P_2) a lower cost if it uses fewer atomic predicates than P_2 (resp. P_1). If P_1 and P_2 use the same number of atomic predicates, then the cost function assigns a lower cost to the program that uses fewer constructs in the column extractors.

Program optimization. While our DSL is designed to facilitate synthesis, programs in this language can be inefficient: In particular, the synthesized programs generate a (potentially large) intermediate table and then filter out the undesired tuples. To avoid inefficiencies caused by this design choice, MITRA-CORE optimizes the synthesized programs by avoiding the generation of intermediate tables whenever possible. In particular, consider a synthesized program of the form $\lambda x. \text{filter}(\pi_1 \times \pi_2, \phi)$. Instead of first taking the cross-product of π_1, π_2 and then filtering out undesired tuples, we optimize the synthesized program in such a way that the optimized program directly generates the final table by using ϕ to guide table generation. More specifically, we use ϕ to find a prefix subprogram π^* that is shared by π_1, π_2 with the property that any subsequent execution of the remaining parts of π_1, π_2 from any node in $\llbracket \pi^* \rrbracket$ yields tuples that are guaranteed to satisfy ϕ . Therefore, the optimized program avoids a post-filtering step and directly generates the output table. A more detailed explanation of this optimization can be found in [57].

Handling full-fledged databases. The synthesis algorithm that we described in Section 5 generates programs for converting a single HDT to a relational table. However, in practice, we would like to use MITRA to convert XML and JSON datasets to a complete database with a given schema. This transformation can be performed by invoking MITRA multiple times (once for each target table) and annotating primary and foreign keys of database tables.

MITRA ensures that the synthesized program obeys primary and foreign key constraints by performing a post-processing step.⁵ To ensure that the primary key uniquely identifies a given row in the table, the synthesized program generates primary keys as follows: If a given row in the database table is constructed from nodes n_1, \dots, n_k in the input tree, then we generate its primary key using an in-

⁵If the primary and foreign keys come from the input dataset, we assume that the dataset already obeys these constraints. Hence, the following discussion assumes that the primary and foreign keys do not appear in the input dataset.

| | Benchmarks | | | Synthesis Time | | Input-output Examples | | | | Synth. Program | |
|----------------|--------------|-----------|-------------|----------------|-------------|-----------------------|-------------|------------|------------|------------------|---------------|
| | #Cols | Total | #Solved | Median | Avg. | #Elements | | #Rows | | #Preds (Avg.) | LOC (Avg.) |
| | | | | (s) | (s) | Median | Avg. | Median | Avg. | | |
| XML | ≤ 2 | 17 | 15 | 0.34 | 0.38 | 12.0 | 15.9 | 3.0 | 4.3 | 1.0 | 13.2 |
| | 3 | 12 | 12 | 0.63 | 3.67 | 19.5 | 47.7 | 4.0 | 3.8 | 2.0 | 17.2 |
| | 4 | 12 | 11 | 1.25 | 3.56 | 16.0 | 20.5 | 2.0 | 2.7 | 3.1 | 19.5 |
| | ≥ 5 | 10 | 10 | 3.48 | 6.80 | 24.0 | 27.2 | 2.5 | 2.6 | 4.1 | 23.3 |
| | Total | 51 | 48 | 0.82 | 3.27 | 16.5 | 27.2 | 3.0 | 3.5 | 2.4 | 17.8 |
| JSON | ≤ 2 | 11 | 11 | 0.12 | 0.27 | 6.0 | 7.4 | 2.0 | 2.7 | 0.9 | 21.3 |
| | 3 | 11 | 11 | 0.48 | 1.13 | 7.0 | 10.5 | 3.0 | 3.5 | 2.0 | 23.0 |
| | 4 | 11 | 11 | 0.26 | 12.10 | 6.0 | 7.9 | 2.0 | 2.8 | 3.0 | 26.5 |
| | ≥ 5 | 14 | 11 | 3.20 | 3.85 | 6.0 | 8.1 | 2.0 | 2.5 | 4.9 | 28.0 |
| | Total | 47 | 44 | 0.31 | 4.33 | 6.0 | 8.5 | 2.0 | 2.9 | 2.7 | 24.7 |
| Overall | 98 | 92 | 0.52 | 3.78 | 11.0 | 18.7 | 3.0 | 3.2 | 2.6 | 21.6 | |

Table 1: Summary of our experimental evaluation

jective function $f(n_1, \dots, n_k)$. Since each row in the table is constructed from a unique list of tree nodes, the generated primary key is guaranteed to be unique for each row as long as f is injective. In our implementation, f simply concatenates the unique identifiers for each tree node.

In order to ensure that a foreign key in table T refers to a primary key in the table T' , the synthesized program for table T must use the same function f to generate the foreign keys. In particular, to generate the foreign key for a given row r constructed from list of tree nodes n_1, \dots, n_k , we need to find the tree nodes n'_1, \dots, n'_m that are used to construct the corresponding row r' in T' . For this purpose, we learn m different (node extractor, node) pairs (χ_j, n_{t_j}) such that $\chi_j(n_{t_j})$ yields n'_j for all rows in the output examples for T and T' . Finally, for a given row r in T , we then generate the foreign key for r as $f(\chi_1(n_{t_1}), \dots, \chi_m(n_{t_m}))$. This strategy ensures that the foreign and primary key constraints are satisfied as long as the learnt node extractors are correct.

7. EVALUATION

To evaluate MITRA, we perform experiments that are designed to answer the following questions:

- Q1.** How effective is MITRA at synthesizing tree-to-table transformation programs?
- Q2.** Can MITRA be used to migrate real-world XML and JSON datasets to the desired relational database?
- Q3.** Are the programs synthesized by MITRA fast enough to automate real-world data transformation tasks?

To answer these questions, we perform two sets of experiments: The first experiment evaluates MITRA on tree-to-table transformation tasks collected from StackOverflow, whereas the second experiment evaluates MITRA on real-world datasets. Both experiments are conducted on a MacBook Pro with 2.6 GHz Intel Core i5 processor and 8 GB of 1600 MHz DDR3 memory running OS X version 10.12.5.

7.1 Accuracy and Running Time

Setup. To perform our first experiment, we collected 98 tree-to-table transformation tasks from StackOverflow by searching for relevant keywords (e.g., “JSON to database”, “XML shredding”, etc.).⁶ Among these 98 benchmarks, 51 involve XML documents, while 47 involve JSON files.

Since MITRA requires input-output examples, we obtained the input XML/JSON file directly from the StackOverflow

⁶All benchmarks are available from [3]

post. For output examples, we used the table provided in the StackOverflow post if one was present; otherwise, we constructed the desired output table based on the English description included in the post.

For each benchmark, we used MITRA to synthesize a program that performs the given task. We manually inspected the tool’s output to check whether the synthesized program performs the desired functionality. Even though any program synthesized by MITRA is guaranteed to satisfy the provided input-output examples, it may not necessarily be the program that the user intended. Whenever the program synthesized by MITRA did not conform to the English description provided in the StackOverflow post, we updated the input-output examples to make them more representative. In our evaluation, we needed to update the original input-output example at most once, and the original examples were sufficient to learn the correct program for the majority of the benchmarks.

Results. Table 1 summarizes the results of evaluating MITRA on these 98 benchmarks. The first part of the table provides information about our benchmarks, which we categorize into four classes depending on the number of columns in the target table. Specifically, “#Cols” shows the number of columns in each category, and “Total” shows the number of benchmarks in each category. The column labeled “#Solved” shows the number of benchmarks that MITRA was able to solve correctly. Overall, MITRA was able to synthesize the target program for 93.9% of these benchmarks.

The columns under “Synthesis Time” show the median and average time that MITRA takes to synthesize the desired program in seconds. On average, MITRA synthesizes the target transformation in 3.8 seconds, and the median time is even lower (0.5 seconds). We believe these results demonstrate that MITRA is quite practical.

Next, the section labeled “Input-output Examples” in Table 1 describes properties of the provided input-output examples. The two columns under “#Elements” represent the median and average number of elements in the input document, and the “#Rows” columns show the median and average number rows in the output table. Here, “#Elements” corresponds to the number of JSON objects and XML elements. As we can see from Table 1, the median number of elements in the input document is 11, and the median number of rows in the input table is 3. These results suggest that MITRA can synthesize the desired program from relatively small input-output examples.

| Name | Dataset | | Database | | Synthesis | | Execution | | |
|----------------|---------|---------|----------|-------|--------------|--------------|-----------|--------------|--------------|
| | Format | Size | #Tables | #Cols | Tot. Time(s) | Avg. Time(s) | #Rows | Tot. Time(s) | Avg. Time(s) |
| DBLP | XML | 1.97 GB | 9 | 39 | 7.41 | 0.82 | 8.312 M | 1166.44 | 129.60 |
| IMDB | JSON | 6.22 GB | 9 | 35 | 33.53 | 3.72 | 51.019 M | 1332.93 | 148.10 |
| MONDIAL | XML | 3.64 MB | 25 | 120 | 62.19 | 2.48 | 27.158 K | 71.84 | 2.87 |
| YELP | JSON | 4.63 GB | 7 | 34 | 14.39 | 2.05 | 10.455 M | 220.28 | 31.46 |

Table 2: Summary of using MITRA for migrating real-world datasets to a full DB. The columns labeled “Tot. Time” include time for all database tables, whereas “Avg. Time” is the average time per table.

The final section of Table 1 describes properties of the synthesized programs. For instance, according to the “Preds” column, the average number of atomic predicates used in predicates ϕ from our DSL is 2.6. More interestingly, the column labeled “LOC” gives the number of lines of code in the synthesized XSLT and Javascript programs. On average, the size of the programs synthesized by MITRA is 21.6 (without including built-in functions, such as the implementation of *getDescendants* or code for parsing the input file).

Limitations. To understand MITRA’s limitations, we investigated the 6 benchmarks for which MITRA was not able to synthesize the desired program. We found that the desired program was not expressible in our DSL for 5 of these 6 benchmarks. For example, some of these 5 benchmarks require the use of a conditional in the column extractor, which our DSL does not currently support. For the other benchmark, there exists a DSL program that can perform the desired functionality, but MITRA ran out of memory due to the large number of columns in the target table.

Performance. We also evaluated the performance of the synthesized programs by running them on XML documents of size 512 ± 20 MB. In particular, we used the Oxygen XML editor [7] to generate XML files with a given schema and a specified number of elements. Among the 48 XSLT programs generated using MITRA, 46 of them were able to perform the desired transformation within approximately one minute.. The running times of these 46 programs range between 8.6 and 65.5 seconds, with a median (resp. average) running time of 20.0 (resp. 23.5) seconds. The other two programs were not able to complete the transformation task within one hour due to inefficiencies in the generated code.

7.2 Migration to Relational Database

Setup. In our next experiment, we use MITRA to convert real-world XML and JSON datasets to a complete relational database. Our benchmarks for this experiment include the following four well-known datasets:

- **DBLP**, an XML document containing 2 GB of data about published computer science papers [2].
- **IMDB**, a set of JSON documents containing 6.2 GB of data about movies, actors, studios etc. [4]⁷
- **MONDIAL**, an XML document containing 3.6 MB of geographical data [6].
- **YELP**, a set of JSON documents containing 4.6 GB of data about businesses and reviews [8].

We obtained the target relational database schema for each of these datasets from [1], and certified that all of

⁷The raw data from [4] is provided in tab-separated-values (TSV) format, so we converted it to JSON format using an existing program [5].

them are normalized. To use MITRA to perform the desired data migration task, we manually constructed small input-output examples. For each table, we provided a single pair of input-output examples, in addition to a list of all primary and foreign keys in the database schema. In particular, the average number of elements in the input examples is 16.6 and the average number of rows in each database table is 2.8. Given a suitable input-output example for each target database table, we then ran MITRA with a time limit of 120 seconds. We then manually inspected the synthesized program and verified its correctness. In this experiment, there was no user interaction; the original examples we provided were sufficient for MITRA to synthesize the desired program. Furthermore, for each target table, we were often able to re-use the input examples (e.g., for IMDB, we used 2 different input examples for the 9 target tables).

Results. The results of this experiment are summarized in Table 2. The columns under “Database” show the number of tables and total number of attributes in the target database. In particular, the number of database tables range between 7 and 25 and the number of columns range from 34 to 120.⁸ Next, the two columns under “Synthesis” show total and average synthesis time in seconds, respectively. Specifically, the average time denotes synthesis time per table, whereas total time aggregates over all database tables. As shown in Table 2, average synthesis time ranges between 0.8 to 3.7 seconds per database table.

The last part of Table 2 provides information about the execution time of the synthesized program and size of the generated database. Specifically, according to the “#Rows” column, the total number of rows in the generated database ranges between 27,000 and 51 million. The next two rows provide statistics about the execution time of the synthesized programs on the original full dataset. For example, the average time to generate a target database table for datasets in the range 2 – 6 GB is between 31 and 148 seconds. These statistics show that the programs generated by MITRA can be used to migrate real-world large datasets.

8. RELATED WORK

This paper is related to a long line of work in databases and programming languages. In what follows, we compare and contrast our approach against existing techniques.

Data Exchange. The problem of converting hierarchically structured documents to a relational format is a form of *data exchange* problem, where the goal is to transform a data instance of a source schema into a data instance of a target schema [20]. Due to the difficulty of manually performing

⁸The Mondial database consists of a large number of tables and columns because it includes a variety of different geographical and cultural data.

such transformations, there has been significant work on automating data exchange tasks [40, 19, 43, 38, 21]. A common approach popularized by *Clio* [40] decomposes the data exchange task into two separate phases: *schema mapping* and *program generation*. A schema mapping describes the relationship between the source and target schemas and is typically specified in the form of declarative logic constraints, such as GLAV (Global-and-Local-As-View) constraints [32]. Given a schema mapping, the second program generation phase “translates” this mapping into executable code [19].

Because manual construction of schema mappings requires non-trivial effort from data architects, there has been significant work on automatically inferring such mappings from various kinds of informal specifications provided by the user. For instance, users may specify element correspondences by drawing lines between elements that contain related data [19, 31, 36, 16, 39, 18]. More recently, examples have become more popular as a way of communication between users and the system. Such example-based approaches can be broadly categorized into two classes depending on what the examples are used for. One approach uses examples to specify the desired schema mapping [41, 11, 9, 12]. To the best of our knowledge, all of these approaches use GLAV (or similar formalisms) as the schema mapping language, and hence can only handle cases where the source and target schemas are both relational. The other approach uses examples to help users understand and refine the generated schema mappings [58, 10], and the learning procedure still takes visual correspondences (lines between elements) as specification. Our work can be viewed as an instantiation of the first approach. However, our method is different from previous methods of this approach in three important aspects. First, we can synthesize programs that convert tree-structured to relational data. Second, we design a DSL as the intermediate language which can express mappings between hierarchical and relational formats. Finally, our PBE-based synthesis algorithm combines finite automata and predicate learning to enable efficient mapping generation.

Programming-by-Example (PBE). The problem of automatically synthesizing programs that satisfy a given set of input-output examples has been the subject of research in the past four decades [46]. Recent advances in algorithmic and logical reasoning techniques have led to the development of PBE systems in a variety of domains including string transformations [25, 47], data filtering [54], data imputation [52], data structure manipulations [23, 56], matrix transformation [53], table transformations [22, 27], SQL queries [51, 60], and map-reduce distributed programs [48].

Among existing PBE techniques, one particularly related work is HADES [56], which can be used to synthesize transformations between tree-structured data, including XML files, from input-output examples. While our internal HDT representation is inspired by HADES, there are some subtle differences: In particular, instead of mapping each element to a single node as in HADES, we map each element to multiple tree nodes. Despite the similarity between the HDT representations, HADES focuses on tree-to-tree transformations and cannot automate transformations from hierarchical to relational data. Although a relational table can be represented as an HDT, HADES’ approach of decomposing trees to a set of paths omits relations between different paths (i.e., columns in the relational table) and can therefore not automate most kinds of interesting tree-to-table transformations.

To the best of our knowledge, the only prior PBE-based tool that aims to automate transformations from semi-structured to relational data is FLASHEXTRACT [34]. While the main focus of FLASHEXTRACT is data stored in spreadsheets, it also supports some tree-structured data, such as HTML documents. However, FLASHEXTRACT is less expressive than MITRA, as it cannot infer relations between different nodes in the tree structure.

XML-to-Relational Mapping. There is a significant body of research on using relational database management systems to store and query XML (and JSON) documents [33, 44, 17, 59, 29, 50, 45, 13, 14]. A typical approach for this problem consists of three steps: First, the tree structure of the document is converted into a flat, relational schema; next, the XML document is “shredded” and loaded into the relational database tables, and, finally, XML queries are translated into corresponding SQL queries. The goal of these systems is quite different from ours: In particular, they aim to efficiently answer queries about the XML document by leveraging RDBMS systems, whereas our goal is to answer SQL queries on a desired relational representation of the underlying data. Furthermore, our XML-to-Relational mapping algorithm is also different from existing XML shredding techniques which can be broadly categorized into two classes, namely structure-centric and schema-centric approaches. The first approach relies on the XML document structure to guide the mapping process [50, 49], whereas the second one makes use of schema information and annotations to derive a mapping [13, 24, 55]. In contrast, our technique uses instances of the input XML document and the output relational table in the mapping process. Moreover, we have designed a novel DSL that can express a rich class of mapping programs and is amenable to efficient synthesis. Finally, the combination of techniques based on DFA and predicate learning in our synthesis algorithm is different from previous XML shredding approaches.

9. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a new PBE-based method for transforming hierarchically structured data, such as XML and JSON documents, to a relational database. Given a small input-output example, our tool, MITRA, can synthesize XSLT and Javascript programs that perform the desired transformation. The key idea underlying our method is to decompose the synthesis task into two phases for learning the column and row extraction logic separately. We have evaluated our method on examples collected from Stackoverflow as well as real-world XML and JSON datasets. Our evaluation shows that MITRA is able to synthesize the desired programs for 94% of the Stackoverflow benchmarks and for all of the real-world datasets.

For future work, we are interested in further optimizing the data migration programs generated by MITRA. In particular, our method performs data migration by running a separate program per database table. However, because these programs operate over the same dataset, we can reduce overall execution time by memoizing results across different programs. Another limitation is that our synthesis algorithm does not return anything if there is no DSL program that satisfies the given examples. To provide more meaningful results to the user in such cases, we plan to investigate techniques that can synthesize programs that maximize the number of satisfied input-output examples.

Acknowledgments

We would like to thank Joanna Bridgwater for her help with implementing MITRA. We also thank members in the UToPiA group and the anonymous reviewers for their thorough and helpful comments. This work was supported in part by NSF Award #1712060, NSF Award #1453386 and AFRL Award #8750-14-2-0270. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government.

10. REFERENCES

- [1] Database schemas. goo.gl/xRMDTe.
- [2] Dblp dataset. <http://dblp.uni-trier.de/xml/>.
- [3] Evaluation benchmarks. <https://goo.gl/JtYXj6>.
- [4] Imdb dataset. <http://www.imdb.com/interfaces>.
- [5] Imdb to json. <https://github.com/oxplot/imdb2json/blob/master/Readme.md>.
- [6] Mondial dataset. <http://www.dbis.informatik.uni-goettingen.de/Mondial/>.
- [7] Oxygen xml editor. <https://www.oxygenxml.com>.
- [8] Yelp dataset. https://www.yelp.com/dataset_challenge.
- [9] B. Alexe, B. T. Cate, P. G. Kolaitis, and W.-C. Tan. Characterizing schema mappings via data examples. *TODS*, 36(4):23, 2011.
- [10] B. Alexe, L. Chiticariu, R. J. Miller, and W.-C. Tan. Muse: Mapping understanding and design by example. In *ICDE*, pages 10–19. IEEE, 2008.
- [11] B. Alexe, B. Ten Cate, P. G. Kolaitis, and W.-C. Tan. Designing and refining schema mappings via data examples. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 133–144. ACM, 2011.
- [12] B. Alexe, B. Ten Cate, P. G. Kolaitis, and W.-C. Tan. EIRENE: Interactive design and refinement of schema mappings via data examples. *VLDB*, 4(12):1414–1417, 2011.
- [13] S. Amer-Yahia, F. Du, and J. Freire. A comprehensive solution to the XML-to-relational mapping problem. In *Proceedings of the 6th annual ACM international workshop on Web information and data management*, pages 31–38. ACM, 2004.
- [14] M. Atay, A. Chebotko, D. Liu, S. Lu, and F. Fotouhi. Efficient schema-based XML-to-Relational data mapping. *Information Systems*, 32(3):458–476, 2007.
- [15] M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, and D. Tarlow. Deepcoder: Learning to write programs. *arXiv preprint arXiv:1611.01989*, 2016.
- [16] H.-H. Do and E. Rahm. COMA: a system for flexible combination of schema matching approaches. In *VLDB*, pages 610–621, 2002.
- [17] I. Dweib, A. Awadi, S. E. F. Elrhman, and J. Lu. Schemaless approach of mapping XML document into Relational Database. In *Computer and Information Technology, 2008. CIT 2008. 8th IEEE International Conference on*, pages 167–172. IEEE, 2008.
- [18] H. Elmeleegy, M. Ouzzani, and A. Elmagarmid. Usage-based schema matching. In *ICDE*, pages 20–29. IEEE, 2008.
- [19] R. Fagin, L. M. Haas, M. Hernández, R. J. Miller, L. Popa, and Y. Velegrakis. Clio: Schema mapping creation and data exchange. In *Conceptual modeling: foundations and applications*, pages 198–236. Springer, 2009.
- [20] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: semantics and query answering. *Theoretical Computer Science*, 336(1):89–124, 2005.
- [21] R. Fagin, P. G. Kolaitis, and L. Popa. Data exchange: getting to the core. *ACM Transactions on Database Systems (TODS)*, 30(1):174–210, 2005.
- [22] Y. Feng, R. Martins, J. Van Geffen, I. Dillig, and S. Chaudhuri. Component-based synthesis of table consolidation and transformation tasks from examples. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2017, pages 422–436. ACM, 2017.
- [23] J. K. Feser, S. Chaudhuri, and I. Dillig. Synthesizing data structure transformations from input-output examples. *SIGPLAN Not.*, 50(6):229–239, June 2015.
- [24] K. Fujimoto, D. D. Kha, M. Yoshikawa, and T. Amagasa. A Mapping Scheme of XML Documents into Relational Databases using Schema-based Path Identifiers. In *International Workshop on Challenges in Web Information Retrieval and Integration*, pages 82–90, 2005.
- [25] S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, pages 317–330. ACM, 2011.
- [26] S. Gulwani. Synthesis from examples: Interaction models and algorithms. In *Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), 2012 14th International Symposium on*, pages 8–14. IEEE, 2012.
- [27] W. R. Harris and S. Gulwani. Spreadsheet table transformations from examples. In *ACM SIGPLAN Notices*, volume 46, pages 317–328. ACM, 2011.
- [28] J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation (3rd Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2006.
- [29] H. Jiang, H. Lu, W. Wang, and J. X. Yu. XParent: An efficient RDBMS-based XML database system. In *ICDE*, pages 335–336. IEEE, 2002.
- [30] Z. Jin, M. R. Anderson, M. Cafarella, and H. V. Jagadish. Foofah: Transforming data by example. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, pages 683–698. ACM, 2017.
- [31] J. Kang and J. F. Naughton. On schema matching with opaque column names and data values. In *SIGMOD*, pages 205–216. ACM, 2003.
- [32] P. G. Kolaitis. Schema mappings, data exchange, and metadata management. In *PODS*, pages 61–75. ACM, 2005.
- [33] R. Krishnamurthy. *Xml-to-sql query translation*. PhD thesis, University of Wisconsin–Madison, 2004.

- [34] V. Le and S. Gulwani. Flashextract: A framework for data extraction by examples. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14*, pages 542–553. ACM, 2014.
- [35] H. Lieberman. *Your wish is my command: Programming by example*. Morgan Kaufmann, 2001.
- [36] J. Madhavan, P. A. Bernstein, and E. Rahm. Generic Schema Matching with Cupid. In *VLDB*, pages 49–58, 2001.
- [37] E. J. McCluskey. Minimization of boolean functions. *Bell Labs Technical Journal*, 35(6):1417–1444, 1956.
- [38] R. J. Miller, L. M. Haas, and M. A. Hernández. Schema mapping as query discovery. In *VLDB*, volume 2000, pages 77–88, 2000.
- [39] A. Nandi and P. A. Bernstein. HAMSTER: using search clicklogs for schema and taxonomy matching. *VLDB*, 2(1):181–192, 2009.
- [40] L. Popa, Y. Velegrakis, M. A. Hernández, R. J. Miller, and R. Fagin. Translating web data. In *Proceedings of the 28th international conference on Very Large Data Bases*, pages 598–609. VLDB Endowment, 2002.
- [41] L. Qian, M. J. Cafarella, and H. Jagadish. Sample-driven schema mapping. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 73–84. ACM, 2012.
- [42] W. V. Quine. The problem of simplifying truth functions. *The American Mathematical Monthly*, 59(8):521–531, 1952.
- [43] M. Roth and W.-C. Tan. Data Integration and Data Exchange: It’s Really About Time. In *CIDR*, 2013.
- [44] J. Shanmugasundaram, E. Shekita, J. Kiernan, R. Krishnamurthy, E. Viglas, J. Naughton, and I. Tatarinov. A general technique for querying xml documents using a relational database system. *ACM SIGMOD Record*, 30(3):20–26, 2001.
- [45] J. Shanmugasundaram, K. Tufte, C. Zhang, G. He, D. J. DeWitt, and J. F. Naughton. Relational databases for querying XML documents: Limitations and opportunities. In *VLDB*, pages 302–314. Morgan Kaufmann Publishers Inc., 1999.
- [46] D. E. Shaw, W. R. Swartout, and C. C. Green. Inferring lisp programs from examples. In *IJCAI*, pages 260–267, 1975.
- [47] R. Singh. Blinkfill: Semi-supervised programming by example for syntactic string transformations. *Proceedings of the VLDB Endowment*, 9(10):816–827, 2016.
- [48] C. Smith and A. Albarghouthi. MapReduce Program Synthesis. PLDI, pages 326–340. ACM, 2016.
- [49] S. Soltan and M. Rahgozar. A clustering-based scheme for labeling xml trees. *International Journal of Computer Science and Network Security*, 6(9A):84–89, 2006.
- [50] I. Tatarinov, S. D. Viglas, K. Beyer, J. Shanmugasundaram, E. Shekita, and C. Zhang. Storing and querying ordered XML using a relational database system. In *SIGMOD*, pages 204–215. ACM, 2002.
- [51] C. Wang, A. Cheung, and R. Bodik. Synthesizing highly expressive sql queries from input-output examples. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 452–466. ACM, 2017.
- [52] X. Wang, I. Dillig, and R. Singh. Synthesis of data completion scripts using finite tree automata. *Proc. ACM Program. Lang.*, 1(OOPSLA):62:1–62:26, Oct. 2017.
- [53] X. Wang, I. Dillig, and R. Singh. Program Synthesis using Abstraction Refinement. In *POPL*. ACM, 2018.
- [54] X. Wang, S. Gulwani, and R. Singh. Fidex: Filtering spreadsheet data using examples. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016*, pages 195–213. ACM, 2016.
- [55] G. Xing, Z. Xia, and D. Ayers. X2R: A System for Managing XML Documents and Key Constraints Using RDBMS. In *ACMSE*, 2007.
- [56] N. Yaghmazadeh, C. Klinger, I. Dillig, and S. Chaudhuri. Synthesizing transformations on hierarchically structured data. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '16*, pages 508–521, New York, NY, USA, 2016. ACM.
- [57] N. Yaghmazadeh, X. Wang, and I. Dillig. Automated migration of hierarchical data to relational tables using programming-by-example. *CoRR*, abs/1711.04001, 2017.
- [58] L. L. Yan, R. J. Miller, L. M. Haas, and R. Fagin. Data-driven understanding and refinement of schema mappings. In *SIGMOD*, volume 30, pages 485–496. ACM, 2001.
- [59] M. Yoshikawa, T. Amagasa, T. Shimura, and S. Uemura. XRel: a path-based approach to storage and retrieval of XML documents using relational databases. *ACM Transactions on Internet Technology*, 1(1):110–141, 2001.
- [60] S. Zhang and Y. Sun. Automatically synthesizing sql queries from input-output examples. In *Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on*, pages 224–234. IEEE, 2013.