

# On Extracting Data from Tables that are Encoded using HTML

Juan C. Roldán\*, Patricia Jiménez, Rafael Corchuelo

*University of Seville, ETSI Informática  
Avda. Reina Mercedes s/n, Sevilla E-41012, Spain*

---

## Abstract

Tables are a common means to display data in human-friendly formats. Many authors have worked on proposals to extract those data back since this has many interesting applications. In this article, we summarise and compare many of the proposals to extract data from tables that are encoded using HTML and have been published between 2000 and 2018. We first present a vocabulary that homogenises the terminology used in this field; next, we use it to summarise the proposals; finally, we compare them side by side. Our analysis highlights several challenges to which no proposal provides a conclusive solution and a few more that have not been addressed sufficiently; simply put, no proposal provides a complete solution to the problem, which seems to suggest that this research field shall keep active in the near future. We have also realised that there is no consensus regarding the datasets and the methods used to evaluate the proposals, which hampers comparing the experimental results.

*Keywords:* HTML documents; web tables; table mining; data extraction.

---

## 1. Introduction

Tables are a common means of displaying data in web documents because people can easily spot and interpret them [3, 5]. The estimations are as high as hundreds of millions; for instance, Lehmborg et al. [35] and Galkin et al. [25] found 233 and 12 billion tables in different editions of the Common Web Crawl, respectively, and Crestan and Pantel [16] found 8.2 billion tables in their own crawl. Cafarella et al. [5] also highlighted the explosion of consumer demand for data that comes from tables thanks to the increasing popularity of voice assistants and infobox-like search results.

In this context, data extraction consists in transforming tables into structured formats that focus on their data and abstract away from how they are

---

\*Corresponding author.

*Email addresses:* jcroldan@us.es (Juan C. Roldán), patriciajimenez@us.es (Patricia Jiménez), corchu@us.es (Rafael Corchuelo)

displayed. Data extraction has many applications to text mining [24, 64, 65], data (meta-)search [3, 9, 18, 26, 44, 51, 63–65], query expansion [16], document summarisation [40, 64], question answering [1, 20, 44, 46, 65], knowledge discovery [9, 22, 26, 32, 44, 46], knowledge base construction [17, 72], knowledge augmentation [1, 9, 18, 20, 56, 56, 57, 67], synonym finding [1, 3, 39], improving accessibility [43, 47, 49, 64, 65], textual advertising [15], data compression [2, 49], or creating linked data [22, 33], just to mention a few common ones.

It is not surprising then that many researchers have worked on a variety of proposals to extract data from tables, which has motivated others to write articles in which they summarise and compare them. Lopresti and Nagy [41, 42] presented a definition of table, with a focus on how they are encoded and displayed, and motivated the need to extract data from them; they summarised some data extraction techniques, as well as some techniques to integrate the resulting data. Hurst [27] introduced the problem and then reported on some of the challenges regarding locating tables and their cells; he paid special attention to reporting on the evaluation of the proposals and concluded that common evaluation methods are not suitable. Zanibbi et al. [70] described the extraction tasks as abstract machine-learning procedures in which input documents are first modelled and then mapped onto observations that are transformed prior to performing inference; they analysed many existing proposals according to how they address the steps of the previous procedure; they also highlighted the need for common evaluation methods. Costa-Silva et al. [14] discussed on what a table is and what makes it different from a diagram; they then listed many proposals to implement the tasks involved in extracting data from tables and compared them using several comparison frameworks; they also criticised common evaluation methods and contributed with some specific purpose evaluation measures. Embley et al. [21] first discussed on the definition of table and then motivated the need to extract data from them by describing many applications; they listed some proposals to locate tables and their cells, but their emphasis was on the tasks to classify the cells, to group them, and to interpret the tables.

The previous articles focus on the proposals that were published between 1990 and 2003. Unfortunately, there is not a recent article that summarises and compares the proposals that were published later, which motivated us to work on it. Our focus is on proposals that work on tables that are encoded using HTML because there has been a steady shift towards encoding them using this language [3, 18, 35], which provides specific-purpose tags and has become pervasive. We have analysed 28 proposals that were published between 2000 and 2018, we have defined a vocabulary that homogenises the terminology used in this field, we have used it to summarise the proposals as homogeneously as possible, and we have compared them side by side using several objective characteristics. We have identified several challenges to which no proposal provides a conclusive solution and also several challenges that have not been addressed sufficiently; addressing them in future shall definitely help produce solutions that increase the range of tables from which data can be extracted correctly. We have also realised that there is not a standardised evaluation method, which hampers the experimental comparison.

The rest of the article is organised as follows: Section 2 introduces the vocabulary that we have compiled; Section 3 summarises the proposals that we have analysed using the previous vocabulary; Section 4 compares them side by side using objective characteristics; finally, Section 5 concludes the article.

## 2. Vocabulary

In this section, we have made a point of integrating the many complementary terms that are commonly used in the literature under a common vocabulary. We first report on the vocabulary that is related to tables themselves and then on the vocabulary that is related to extracting data from them. We illustrate most of the concepts with a couple of examples.

### 2.1. Table-related vocabulary

Unfortunately, there is not a consensus definition in the literature regarding what a table is. Many authors focus on the encoding since they define them as whatever one can encode within HTML table tags [3, 7, 9, 16, 18, 28, 30, 32, 34, 38, 47, 49, 59, 66, 68], which is a pragmatic approach; a few also refer to the display of data, since they define tables as grids in which data are located in cells in a manner that lines and/or styles ease interpreting them [22, 24, 26, 28, 30, 32, 46, 49]. There is only a proposal that deviates a little from the previous approaches [21] since the authors focus on the data model behind the tables, independently from how they are displayed; their proposal, however, works on tables in which data are arranged in grids.

Neither is there a consensus taxonomy of tables. Most authors differentiate between data tables, which provide data to be extracted, and non-data tables, which are used for layout purposes or to provide utilities. Many of them make also a difference between listings, forms, matrices, and enumerations [16, 18, 26, 30, 34, 44, 46, 66], although the exact terminology used is very diverging; there is also a proposal in which tables are classified according to whether they have headers or not [22].

In the previous discussion, there are three key concepts, namely: encoding, cell, and table, which we define below.

**Definition 1 (Encodings):** An encoding is a specification of how a table must be displayed to a person. Common encodings include pre-formatted text, images, and mark-ups. In a table that is encoded using pre-formatted text, the data are arranged in lines, they are aligned to their corresponding columns using blanks, and the cells may be delimited using, for instance, dashes, vertical bars, or tabulators. In a table that is encoded using an image, there is a graphic canvas onto which the data and the lines that delimit the cells, if any, are drawn using bitmaps or vectors. Contrarily, mark-ups provide a variety of tags that help encode the tables, their cells and, hopefully, additional information that helps interpret them. There are several mark-up languages available [58], but our focus is on HTML due to its pervasiveness in the Web. HTML provides an array of table-related tags, namely: `table`, `thead`, `tbody`, `tfoot`, `col`, `colgroup`,

th, tr, td, and caption. It is relatively easy to extract data from tables that are encoded using the previous tags. Unfortunately, real-world tables have a variety of intricacies that hamper the extraction process, namely: some tables are encoded using a subset of table-related tags that hardly help locate them and their cells, which does not help interpret them; other tables are encoded using listing tags (ul, ol, dl, li, dd, and dt) [9, 20, 36, 37]; lately, it is also relatively common to find tables that are encoded using block tags (div and span) due to their ability to create responsive layouts [50]; and, generally, speaking, there are many tables that are encoded using a variety of tags that are not actually related to tables, but look like tables when they are displayed [24, 26].  $\square$

**Definition 2 (Cells):** A cell is a box that provides contents to a table. They can be classified along several axes, namely: a) According to how they are segmented, cells can be single cells, which occupy exactly one position in the grid of a table, or spanned cells, which occupy more than one position. b) According to whether their contents are complete or not, cells can be classified as single-part cells, whose contents are complete, and multi-part cells, which provide partial contents that must be somewhat merged with the contents of other cells. c) According to their function, they can be classified as meta-data cells, whose contents are labels that help people understand other contents in the table, data cells, whose contents provide the data that must be extracted, decorator cells, which provide irrelevant contents, and context-data cells, which provide captions, notes, or factorised data. d) According to how their contents must be interpreted, cells can be classified as factorised cells, whose contents must be borrowed from adjacent cells, void cells, which are not intended to provide any contents, atomic cells, whose contents cannot be decomposed further, and structured cells, whose contents can be decomposed into a mixture of data and meta-data.  $\square$

**Definition 3 (Tables):** A table is a collection of cells that are arranged in rows and columns within a grid, where lines and/or styles are typically used to help people interpret them. There are cases in which some context data are provided in the text that surrounds a table, i.e., captions, notes, and factorised data. The cells in a table are typically grouped as follows according to their functions: headers, which are groups of meta-data cells, tuples, which are groups of data cells, and separators, which are groups of decorator cells. Typically, headers are arranged on the first few rows and/or columns, but we have found some tables in which they are interwoven with tuples for the sake of readability; it was the case of long listings with many tuples, in which it makes sense to repeat the header rows or columns every few tuples, or wide listings/forms with many headers, in which it makes sense to split the header rows or columns to narrow them. Data tables can be broadly classified as follows: a) listings, in which the headers, if any, occupy either the first few rows or columns and the tuples are arranged in the remaining rows or columns, respectively; b) forms, in which the headers, if any, occupy either the first few rows or columns and there is a single tuple that is arranged row- or column-wise, respectively; c) matrices, in which the headers occupy both the first few rows and columns, and all of the

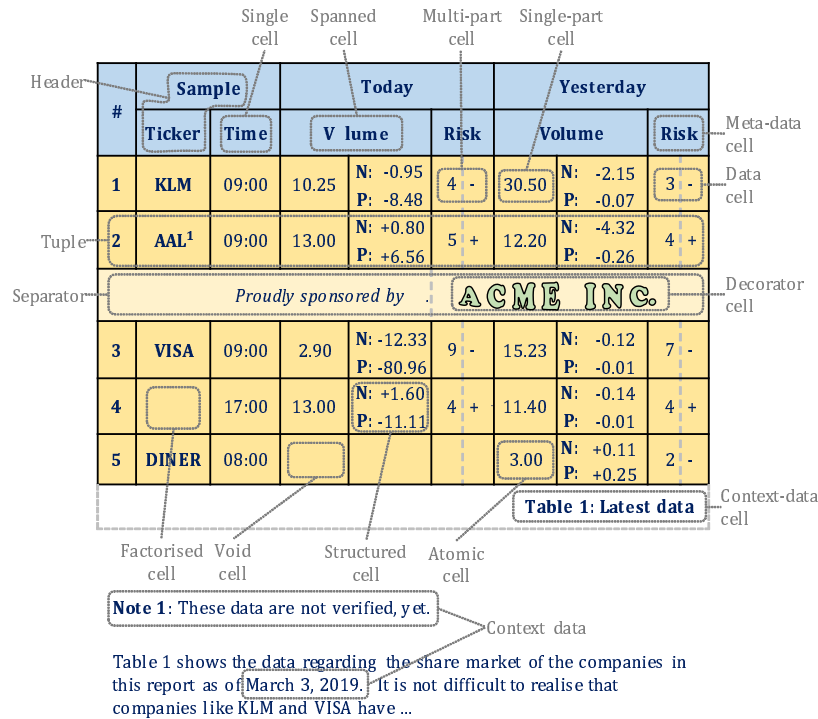


Figure 1: A sample table.

data cells constitute a single tuple; and d) enumerations, in which there are no headers and each individual cell can be considered a tuple. According to Crestan and Pantel [16], this taxonomy covers roughly 98% of the data tables in their 8.2-billion table repository; the authors mention that it is arguable that the remaining 2% tables can be considered actual data tables or that they are frequent enough to be representative. □

**Example 1:** Figure 1 shows a horizontal listing taken from a document on a share market. The black, solid lines help delimit the boundaries of the cells in the grid; the greyed, dashed lines represent the boundaries of a few cells that exists in the encoding of the table, but are not visible to the reader because they are used for layout purposes. Most of the data are displayed on an  $9 \times 11$  grid, but there are also some context data in the surrounding text.

Cells like “Time” are single because they occupy exactly one position in the grid; on the contrary, cell “Volume” is spanned because it occupies two positions in the grid. The meta-data cells occupy the first two rows; for instance “Risk” is one such meta-data cell. Contrarily, cell “3-” below is a data cell and cell “Acme Inc.” is a decorator cell. The caption of the table is displayed within a bottom cell that spans the whole table, which is considered a context-data cell; realise that there are additional context data: there is a note regarding cell

“AAL” and there is a factorised datum regarding the date of the report, which complements the times provided in some cells. Cells like “30.50” are single-part cells because their contents are complete; on the contrary, cells like “4-” are multi-part cells because it is necessary to merge the contents of two cells so that the contents of the resulting cell are complete. A cell like “3.00” is atomic since its contents cannot be decomposed further; contrarily, a cell like “N: +1.60 P: -11.04” is a structured cell because it provides both meta-data and data, which means that it can be decomposed further. The empty cell below cell “VISA” very likely factorises the ticker since there are two tuples regarding this company at 09:00 and 17:00, respectively; contrarily, the four empty cells on the right of cell “08:00” are very likely void cells that indicate that no data are available.

The table has seven headers, namely: “#”, “Sample/Ticker”, “Sample/Time”, “Today/Volume”, “Today/Risk”, “Yesterday/Volume”, and “Yesterday/Risk”. It provides five tuples, the first of which is (“1”, “KML”, “09:00”, “10.25”, “N: -0.95 P: -8.48”, “4-”, “30.50”, “N: -2.15 P: -0.07”, “3-”). It has also a separator at the fifth row, which shows an advertisement.

## 2.2. Data-extraction vocabulary

In our context, data extraction refers to a process that transforms the tables in an input document into record sets. A record is a data structure in which the individual data in a tuple are endowed with semantics by means of descriptors that are computed from the meta-data provided by the corresponding table; in cases in which the table does not provide enough meta-data, the descriptors must be generated artificially.

Costa-Silva et al. [14] did a good job at identifying the tasks of which the data-extraction process is composed, namely: location, segmentation, functional analysis, structural analysis, and interpretation. Note, however, that their focus was on tables that are encoded using pre-formatted text or images, which means that they need not make tables that provide data apart from tables that are intended for layout purposes or to provide utilities. The latter are very common in nowadays Web, which motivated Cafarella et al. [3], for instance, to introduce a task to discriminate data tables from non-data tables.

Before feeding the record sets returned by data extraction into a particular application, it is commonly necessary to perform some of the following integration tasks: semantisation [25, 45, 54, 55, 60, 63, 71], which either maps the descriptors onto the terminology box of a particular ontology or the tuples onto its assertion box [19]; union [23], which merges record sets that provide similar data; finding primary keys [62], which determines which components of the tuples identify them as univocally as possible; record linkage [8, 11, 12], which finds different records that refer to the same actual entities; augmentation [6, 52, 67], which joins record sets on the same topic to complete the information that they provide individually; and cleaning [10, 31, 61], which fixes data. Note that the integration tasks are orthogonal to data extraction because they are independent from the source of the record sets, which is the reason why they fall out of the scope of this article.

In the previous discussion, there are three key concepts: record set, extraction task, and data extraction, which we define below.

**Definition 4 (Record set):** A record set is a collection of records. A record is a map that associates a set of descriptors to each of the components of a tuple. A descriptor is a structured label that endows the components of a tuple with the semantics provided by the meta-data in the corresponding headers or structured cells; if not enough meta-data are available, then descriptors must be generated artificially. We make three types of descriptors apart, namely: simple descriptors, which correspond to the contents of a single meta-data cell, field descriptors, which correspond to the contents of several adjacent meta-data cells, and artificial descriptors, which are used when not enough meta-data are available. In listings and forms, every component of the tuples has one associated descriptor; in matrices, they have two associated descriptors; in enumerations, the descriptors must be created from the meta-data in the cells, if any; in other cases, they must be generated artificially.  $\square$

**Definition 5 (Extraction tasks):** The tasks involved in extracting data from a table are the following [14]: a) location, which searches the input document for the excerpts in which tables are encoded and returns them; b) segmentation, which searches for the cells of which a table is composed; c) discrimination, which classifies a table as either a data table or a non-data table, but further sub-classification is possible; d) functional analysis, which classifies the cells according to their functions; e) structural analysis, which groups cells into at least headers and tuples; and f) interpretation, which produces record sets building on the results of the previous tasks.  $\square$

**Definition 6 (Data extraction):** Data extraction refers to a process that organises the extraction tasks into a pipeline so that they can achieve their goal. Zanibbi et al. [70] and Costa-Silva et al. [14] reported on the many common inter-dependencies amongst the extraction tasks.  $\square$

**Example 2:** Figure 2 illustrates a sample data extraction process in which we have organised the tasks into a sequential pipeline.

The location task finds two excerpts in the input document that seem to have tables; the segmentation task is responsible for finding the individual cells of which the tables are composed, plus the context data that is associated with them; the discrimination task makes a difference between the table on the left, which seems to be a menu that does not provide any data, and the table on the right, which seems to be a table that provides data; the functional analysis task makes meta-data cells apart from data cells; the structural analysis task groups the meta-data cells into four headers and the data cells into two tuples; finally, the interpretation task produces a record set with three records.

Regarding the descriptors, we illustrate them using the usual field-access notation for simple and field descriptors and the usual array-access notation for artificial descriptors. For instance, header “A/A” results in a simple descriptor of the form “A” because both cells were actually a vertically-spanned cell in the original table. On the contrary, header “B/C” results in a field descriptor of the form “B.C” in which it is clear that whatever “C” represents is subordinated to




Figure 2: A sample data extraction process.

whatever “B” represents; note that this descriptor is ambiguous since there are two columns of the table with the same header. In such cases, the table does not provide enough meta-data and the columns must be made apart by means of artificial descriptors, that is “B.C[1]” and “B.C[2]”. Obviously, header “B/D” results in a field descriptor of the form “B.D”.

The records extracted are the following: {"A": "e", "B.C[1]": "f", "B.C[2]": "g", "B.D": "h"}, {"A": "i", "B.C[1]": "j", "B.C[2]": "k", "B.D": "l"}, and {"\$caption": "Table 1: XXX"}. Realise that the last record uses a special simple descriptor to indicate that corresponding datum is the caption of the table.



Reference	Location	Segmentation	Discrimination	Functional analysis	Structural analysis	Interpretation
Chen et al. 2000	Naive	Naive	Yes	Yes	No	Naive
Lerman et al. 2001	Yes	Yes	No	No	Yes	No
Penn et al. 2001	Naive	No	Yes	No	No	No
Yoshida and Torisawa 2001	No	No	No	Yes	Naive	No
Cohen et al. 2002	Naive	Yes	Yes	No	Yes	No
Hurst 2002	Naive	No	Yes	No	No	No
Wang and Hu 2002	Naive	No	Yes	No	No	No
Yang and Luk 2002	Naive	Naive	Yes	Yes	Yes	Naive
Lerman et al. 2004	Yes	Yes	No	No	Yes	No
Kim and Lee 2005	Naive	No	Yes	Yes	No	No
Jung and Kwon 2006	Naive	Naive	Yes	Yes	No	No
Gatterbauer et al. 2007	Yes	Yes	Yes	Naive	No	No
Okada and Miura 2007	Naive	No	Yes	No	No	No
Cafarella et al. 2008	Naive	No	Yes	Yes	No	Yes
Crestan and Pantel 2011	Naive	No	Yes	No	No	No
Elmeleegy et al. 2011	No	Yes	No	No	Naive	No
Fumarola et al. 2011	Yes	No	Yes	No	Yes	No
Lautert et al. 2013	No	No	Yes	No	No	No
Ling et al. 2013	Yes	Yes	No	Naive	Naive	No
Son and Park 2013	Naive	No	Yes	No	No	No
Braunschweig et al. 2015	No	No	No	Naive	Naive	No
Chu et al. 2015	Naive	Yes	No	No	Naive	No
Eberius et al. 2015	Naive	No	Yes	No	No	No
Embley et al. 2016	No	Yes	No	Yes	No	Naive
Milošević et al. 2016	Naive	No	No	Yes	Yes	Naive
Wu et al. 2016	Naive	No	Yes	Naive	Naive	Yes
Nishida et al. 2017	No	No	Yes	No	No	No
Liao et al. 2018	Naive	No	Yes	No	No	No

Table 1: Tasks addressed by each proposal.

### 3. Summary of proposals

In this section, we summarise many proposals that have not been surveyed previously; cf. Table 1. Note that our summaries are intended to provide an overall picture using the vocabulary that we have presented previously; please, refer to the original articles for the full descriptions.

#### 3.1. Location

Yoshida et al. [69], Elmeleegy et al. [20], Lautert et al. [34], Braunschweig et al. [1], Embley et al. [22], and Nishida et al. [46] did not pay attention to the location task. Chen et al. [7], Cohen et al. [13], Hurst [28], Yang and Luk [68], Kim and Lee [32], Jung and Kwon [30], Okada and Miura [47], Cafarella

et al. [3], Son and Park [59], Chu et al. [9], Eberius et al. [18], Milošević et al. [44], Wu et al. [66], and Liao et al. [38] reported on naive approaches that consisted in extracting every HTML excerpt with a `table` tag; Penn et al. [49], Wang and Hu [64], and Crestan and Pantel [16] followed the same approach but discarded tables with nested tables. The other proposals provide more sophisticated approaches.

Lerman et al. [36, 37] focused on tables that are encoded using listing tags. Their proposal works as follows: a) first, the input documents are tokenised and the tokens are assigned to lexical types; b) then, the smallest input document is taken as a base template; c) the remaining documents are then iteratively compared to the template in order to make the sequences of tokens that appear exactly once apart from the others; d) finally, the excerpts of the document that have the largest repetitive sequence of tokens are returned. The authors did not evaluate their procedure in isolation, but their complete system.

Gatterbauer et al. [26] presented a visual approach that analyses the bounding boxes used to display the elements of a document in an attempt to identify tables, lists, and so-called aligned graphics that represent tabular structures. Their proposal works as follows: a) first, they apply some heuristics to locate elements in the DOM tree that are likely to be a part of a table, a list, or an aligned graphic; the authors mention that they have compiled a collection of over twenty such heuristics, but they document only twelve of them in their paper; b) then, they apply an algorithm that searches for so-called frames, which are collections of elements that are rendered so that they form a box; c) then, the frames are expanded to four orthogonal directions by finding elements whose bounding boxes are near to each other; d) finally, the excerpts that correspond to the extended frames are returned. The authors evaluated their proposal on 493 tables from their own repository plus 19 additional tables from Wang and Hu’s [64] repository.

Fumarola et al. [24] also presented a visual approach. Their proposal works as follows: a) it first creates a bounding box that encloses the whole input document; b) it then iterates recursively and creates a bounding box for every element in that document; c) next, it analyses the positions of the inner bounding boxes and finds those that are laid out in a row- or a column-wise manner; d) then, the corresponding excerpts are returned. The authors did not evaluate their procedure in isolation, but their complete system.

Ling et al. [39] presented a proposal whose focus is on locating context-data cells. It works as follows: a) first, it locates the elements in the input document that have a `table` tag; b) then, it extracts some context data from the title tag; c) next, it segments the text around the tables and aligns the resulting segments using a multiple string alignment algorithm; d) finally, the segments that are repetitive enough are considered context-data cells. The authors did not evaluate their procedure in isolation, but their whole system.

### 3.2. Segmentation

Penn et al. [49], Yoshida et al. [69], Hurst [28], Wang and Hu [64], Kim and Lee [32], Okada and Miura [47], Cafarella et al. [3], Crestan and Pantel [16],

Fumarola et al. [24], Lautert et al. [34], Son and Park [59], Braunschweig et al. [1], Eberius et al. [18], Milošević et al. [44], Wu et al. [66], Nishida et al. [46], and Liao et al. [38] did not report on any proposals to implement this task. Chen et al. [7], Yang and Luk [68], Jung and Kwon [30] relied on a naive approach that searches for the cells using specific tags. The other proposals provide more sophisticated approaches.

Lerman et al. [36] focused on tables that are encoded using listing tags; implicitly, they assumed that tuples are shown in a row-wise manner. Prior to segmentation, the authors applied a document alignment method to detect the template of the documents and their repetitive segments, which are very likely to contain the lists. Once the lists are located, their proposal works as follows: a) first, the segments are grouped according to their separators; b) then, DataPro is invoked on the previous groups to learn patterns that characterise their data; c) then, for each segment in each group, it computes binary features that indicate whether it matches the previous patterns or not; d) next, the Auto-Class clustering algorithm is invoked to learn the optimal number of clusters and to learn a set of rules that assign new segments to the most similar cluster; e) finally, the data in each cluster is assumed to be a column of the corresponding table, which facilitates identifying the cells in a row-wise manner. The evaluation was performed on the tables from a repository with 50 documents that were taken from 14 different sources. The authors did not evaluate their procedure in isolation, but their complete system.

Cohen et al. [13] relied on some transformations that help normalise tables before they are segmented. Their proposal works as follows: a) the HTML structure is cleaned up using HTML Tidy and the extra cells generated by this tool are removed; b) structured cells are divided into multiple atomic cells by splitting inner tables, paragraphs, or pre-formatted text; c) spanned cells are split into several cells unless this results in more cells than the height or the width of the table. The authors did not evaluate their procedure in isolation, but their complete system.

Lerman et al. [37] segmented tables by learning a probabilistic model from the repetitive segments in which they decompose tables that are encoded using listing tags; they assumed that tuples are shown in a row-wise manner. Their proposal works as follows: a) lists are split into columns according to candidate separators, which can be tags or punctuation symbols; b) some content features are then computed on each column and their siblings; c) then, an inference algorithm learns a probabilistic model from the previous features; d) the parameters are then used to find the best column assignment for a segment, which is the one that maximises the probability of the features observed given the model. Their evaluation was performed on the tables from a repository with 283 tables from 12 web sites on book sellers, property taxes, white reports, and corrections. They also experimented with a constrain satisfaction approach that was less accurate.

Gatterbauer et al. [26] presented a proposal that requires to identify the spatial relationships between the individual cells of a table. It works as follows: a) it computes the boxes that represent the elements of an input document,

taking into account their contents area, padding, border, and margin areas according to the CSS2 visual formatting model; b) it then overlaps a grid that helps identify each box by means of the co-ordinates of its upper-left corner and its lower-right corner; c) then, it aligns the boxes according to their horizontal and vertical projections; d) next, an adjacency relation is computed according to how distant the cells are; e) finally, some cells are selected and a recurrent expansion algorithm is invoked in an attempt to explore the adjacency relation to find their neighbours. The authors evaluated their proposal on the tables provided by a repository with 1 537 documents that were retrieved from search engines, from Wang and Hu’s [64] repository, or written by the authors. The authors did not evaluate their procedure in isolation, but their complete system.

Elmeleegy et al. [20] tried to find columns by checking how similar the cells in a table are. The similarity is analysed by means of their data types and delimiters. The authors used two resources, a large-scale language model, which helps know sentences that should not be split because they have previously occurred within a cell, and a corpus of tables, which helps identify data that appear in the same column in other tables. Their proposal works as follows: a) each row is split into a (possibly different) number of columns using two scoring functions, namely: a field quality score, which measures the quality of an individual column candidate, and a field-to-field consistency score, which measures the likelihood that two column candidates are actually the same column; b) then, it sets the number of columns to the most frequent one; c) padding columns are added to rows that have less columns than expected and some columns are merged otherwise; d) finally, the segmentation of cells is refined by checking the consistency amongst the cells in a per-column basis; e) if the consistency check fails, the procedure is re-launched. Their evaluation was performed on 20 tables from 20 different domains plus 100 additional tables that were randomly sampled from the Web.

Ling et al. [39] assumed that tables can be segmented building on their td tags; their key contribution was regarding how to find context data. Their proposal works as follows: a) it first uses a number of heuristics to generate candidate context data, namely: tokens in between some punctuation marks, the longest common sub-sequences, pieces of text that can be wikified [53], and pieces of text that vary from document to document but are located at the same position; b) then, the previous context data are added to the original table as additional columns; c) finally, a pairwise adaptation of the Multiple Sequence Alignment algorithm is used to segment the context data. The evaluation was performed on 20 000 tables that were picked from a repository with 130 million tables from 10 different web sites.

Chu et al. [9] also focused on finding the columns of a data table. Their proposal works as follows: a) each row is tokenised using a set of user-defined delimiters; b) then candidate columns are generated using two approaches: a seed tuple is provided and the system discards segmentations that are very different; a custom pruning procedure that borrows some ideas from the well-known A\* procedure is also used; c) it then measures the similarity of each column using lexical and semantic similarity functions that are averaged (the

former computes the difference regarding the number of tokens, characters, and pattern-based types; the latter computes the point-wise mutual information function); d) the process is repeated until a segmentation that maximises similarity is found. Their evaluation was performed on 100 million data tables that were transformed into lists; they used 20 additional tables encoded as lists from five different domains.

Embley et al. [22] presented a proposal that works as follows: a) the input documents are transformed into a representation that preserves the contents only; b) spanned cells are split and their contents are copied verbatim to the resulting cells; c) then, every row with more than two empty cells is considered to provide context data; d) finally, the right-most bottom non-empty cell is considered to be the last cell in the table. (Note that they can work on tables that come from spreadsheets, in which it is not uncommon to find empty cells that are not actually part of any tables.) The authors did not evaluate their procedure in isolation, but their complete system.

### 3.3. Discrimination

Lerman et al. [36], Yoshida et al. [69], Lerman et al. [37], Elmeleegy et al. [20], Ling et al. [39], Braunschweig et al. [1], Chu et al. [9], Embley et al. [22], and Milošević et al. [44] did not pay attention to the discrimination task. The other proposals provide sophisticated approaches.

Chen et al. [7] devised a proposal to discriminate tables by means of heuristics. It works as follows: a) a cell similarity measure is computed by combining string similarity, named entity similarity, and number similarity functions; b) then, the tables whose cells do not exceed a threshold regarding the number of similar neighbour cells are discarded; c) finally, tables with less than two cells or tables with many links, forms, or figures, are also discarded. The evaluation was performed on 3 218 tables from their own repository with documents on airlines from the Chinese Yahoo! site.

Penn et al. [49] also devised a heuristic-based approach. Their proposal works as follows: a) tables that do not have multiple rows and columns are discarded; b) tables whose cells have more than one non-text-formatting tag are also discarded; c) finally, tables whose cells have more than a user-defined number of words are also discarded. The authors also mentioned that a desirable feature is to have syntactic and semantic similarity into account, but they did not explore this idea. They experimented with an unspecified number of tables from their own repository with documents from 75 sites on news, television, radio, and companies.

Cohen et al. [13] devised a proposal that builds on machine learning a classifier. It works as follows: a) some structural and content features are computed from a learning set with tables that are pre-classified as either data tables or non-data tables; b) then, several classifiers are machine-learned and evaluated; c) the classifier that achieves the best effectiveness is selected to implement the discrimination task. The authors experimented with Multinomial Naive Bayes, Maximum Entropy, Winnow, and a decision tree learner that was based on C4.5; their conclusion was that the best results were achieved using Winnow.

They evaluated their proposal using a 5-trial approach on 339 tables from their own repository; in each trial, 75% of the tables were used for learning and the remaining 25% for evaluation purposes.

Hurst [28] presented another machine-learning approach in which he also took visual features into account. He performed his evaluation on 89 data tables and 250 non-data tables from his own repository; they were randomly grouped into five sets from which 25% of the tables were selected for learning purposes and 75% for evaluation purposes. The results confirmed that Naive Bayes achieved the best results when the whole set of features was used, whereas Winnow worked better when only geometric features were used.

Wang and Hu [64] devised another machine-learning proposal that relies on structural and content features that are used to feed a custom decision tree learner; some of the features need to be transformed into real values using Naive Bayes or  $k$ -NN. The content features rely on the words found in the input documents, which requires a large learning set so as to minimise the chances that a classifier is applied to a document with a word that was not in the learning set. The evaluation was performed using 9-fold cross evaluation on 11 477 tables from their own repository with documents from Google’s directories.

Yang and Luk [68] reported on another heuristic-based method. Their proposal works as follows: a) tables that have th tags are considered data tables; b) tables that do not only contain links, forms, or images are also considered data tables; c) meta-data and data cells are then located using some user-defined patterns; d) tables that do not have both meta-data and data cells are discarded. They evaluated their method on 1927 tables from their own repository, which was assembled with random documents from the Web.

Kim and Lee [32] used heuristics and an algorithm to check how similar the cells are. Their proposal works as follows: a) tables are considered data tables if they contain caption or th tags and there are td tags at the right or the bottom sides; b) they are discarded if they have a single cell, if they have nested tables, or if they seem to have meta-data cells only; c) if they have too many links, images, or empty cells, then they are also discarded; d) then, it checks that the cells selected previously are consistent using some user-defined patterns; e) if the degree of similarity per row or column does not exceed a pre-defined threshold, then the corresponding table is discarded. The evaluation was performed on 11 477 tables from Wang and Hu’s [64] repository.

Jung and Kwon [30] presented a machine-learning proposal. It works as follows: a) it first removes empty rows and columns, splits spanned cells by duplicating their contents, and discards tables with only one cell; b) then, it computes many structural, visual, and content features of the table to find out if it has meta-data cells, in which case the table is assumed to have data; c) finally, a C4.5 learner is fed with the input features and the classified tables. The evaluation was performed using 10-fold cross evaluation on 10 000 tables from their own repository plus some tables from Wang and Hu’s [64] repository.

Gatterbauer et al. [26] reported on an approach that identifies tables using some display heuristics. Their proposal works as follows: a) elements with td, th, and div tags are considered candidate tables; b) it tries to identify frames

that rely on those elements, which are assumed to be tables; c) overlapping tables are discarded; d) tables are also discarded if, after removing separator columns and rows, they have less than three rows, a single cell is more than 40% the total size of the table, or they contain cells with more than 20 words. The evaluation was performed on 493 tables from their own repository.

Okada and Miura [47] devised another machine-learning approach that requires to binarise discrete features before feeding them into an ID3 learner. The evaluation was performed using 10-fold cross evaluation on 100 data tables and 100 non-data tables from their own repository.

Cafarella et al. [3] proposed another machine-learning approach. Their proposal works as follows: a) it considers tables that have at least four cells, are not embedded in HTML forms, and are not calendars; b) the tables that meet the previous criteria are classified as either data or non-data tables by a person; c) then, a statistical classifier is machine-learned from a dataset that vectorises the previous tables using both structural and content features that are intended to measure how consistent the cells are. They evaluated their proposal using 5-fold cross evaluation over several thousand tables from their own repository.

Crestan and Pantel [16] also presented a machine-learning proposal. It works as follows: a) tables that have less than four cells or have cells with more than 100 characters are discarded; b) next, some structural and content features are computed; c) then, a Gradient Boosted Decision Tree classification model is machine-learned. The evaluation was conducted on 5 000 tables from their own repository by performing 20-fold cross evaluation without overlapping.

Fumarola et al. [24] proposed a heuristic-based approach. Their proposal works as follows: a) it groups the elements whose bounding boxes are arranged in a grid; b) it then computes their similarity by comparing their DOM trees; c) next, it computes the number of nodes in each group; d) if the similarity in a group is above a user-defined threshold and the difference in the number of nodes is below another user-defined threshold, then it is considered a data table. The evaluation was performed on 224 tables that were gathered from Gatterbauer et al.'s [26] repository.

Lautert et al. [34] devised a machine-learning proposal that builds on neural networks. It works as follows: a) it computes some structural, visual, and content features; b) then, it uses them to machine-learn a perceptron with one hidden layer and resilient propagation; c) it has one output neuron per type of data table, which is encoded using a score in range [0.00 .. 1.00]; the classification is performed in two steps, namely: the first one uses 25 features to classify the tables into the corresponding types and the second step uses the previous 25 features plus the type of table output by the previous classifier. The evaluation was performed on a repository with 342 795 tables that were gathered randomly.

Son and Park [59] also tried a machine-learning approach. Their proposal works as follows: a) it selects every DOM node with tag table and their corresponding parents; b) the features described by Wang and Hu [64] are then computed to create a learning set; c) finally, an SVM classifier is machine-learned using a kernel that works with structural features plus a kernel that works with content features; the structure kernel is based on two other kernels, one

of which works on the table nodes and the other on the corresponding parent nodes. The authors performed 10-fold cross evaluation on a subset of 11 477 tables from Wang and Hu’s [64] repository; roughly 89% of the tables were used for learning purposes and roughly 11% were used for evaluation purposes.

Eberius et al. [18] devised a proposal that builds on machine learning a classifier. It works as follows: a) some heuristics are applied to filter most non-data tables out, namely: tables with less than two rows or columns, tables with an invalid HTML structure, and tables that cannot be displayed correctly; b) some structural and content features are then computed regarding the tables and some of their subregions in order to compute local features; c) two alternatives are now tried: learning one classifier for every table type or using one classifier to discriminate between data and non-data tables and an additional classifier to classify some kinds of data tables; d) several classifiers are machine-learned and evaluated, namely: CART, C4.5, SVM, and Random Forest; e) the classifier that achieves the best effectiveness is selected to implement the discrimination task. They evaluated their proposal on a repository with 24 654 tables from the October 2014 Common Crawl. According to their experience, the best results were achieved with Random Forest.

Wu et al. [66] provided a method to cluster tables that are similar according to their structure. Their proposal works as follows: a) for every two tables, it computes the set of paths that corresponds to `caption`, `td`, and `th` tags; b) then, the similarity between the paths of every two tables is computed; c) then, tables are clustered according to their local density plus the previous similarities; d) now, for each cluster, clustering is performed again building on the paths that lead to elements with tags `li`, `span`, or `div`; e) finally, a so-called artificial judgment method is used to decide on the class of each cluster. The authors used a repository with 5 000 tables from the Wikipedia to evaluate their system, but no results were provided regarding this task.

Nishida et al. [46] devised a proposal that analyses a subset of cells at the top-left corner of a table using a deep neural network. It works as follows: a) for each `td` or `th` tag, an embedding is generated by tokenising words, tags, and row and column indexes; b) each token is encoded as a one-hot vector; c) an LSTM with an attention mechanism is then used to obtain a semantic representation of each cell; d) a convolutional neural network is then connected to three residual units and applied to vectorise the input table; e) finally, a classification layer is used. The authors learnt the network using 3 567 tables from 200 web sites, and evaluated the results on 60 678 tables from 300 web sites; the documents were selected from the April 2016 Common Crawl. They also experimented with an ensemble of five neural networks, which attained the best results.

Liao et al. [38] presented a heuristic-based approach that takes into account the existence of nested data tables. It works as follows: a) tables with a `th` or `caption` tag are considered data tables; b) tables with a large number of pictures, frames, forms, or script tags are discarded; c) tables with a small number of elements or many empty cells are discarded, too; d) tables with too many homogeneous contents in their rows are considered incomplete data tables, which must be stitched to other sibling tables to create a complete data table.



They evaluated their method on 226 tables from 50 different sites.

#### 3.4. Functional analysis

Lerman et al. [36], Penn et al. [49], Cohen et al. [13], Hurst [28], Wang and Hu [64], Lerman et al. [37], Okada and Miura [47], Crestan and Pantel [16], Elmeleegy et al. [20], Fumarola et al. [24], Lautert et al. [34], Son and Park [59], Chu et al. [9], Eberius et al. [18], Nishida et al. [46], and Liao et al. [38] did not report on any proposals to implement the functional analysis task. Gatterbauer et al. [26] presented a naive approach that matches the structure of a table to a number of pre-defined structures in which it is also relatively easy to find the meta-data cells. Ling et al. [39] and Wu et al. [66] assumed that meta-data cells can be easily located by searching for `th` tags. Braunschweig et al. [1] also presented a naive solution since they assumed that meta-data cells are located on the first row. The other proposals provide more sophisticated approaches.

Chen et al. [7] devised a proposal that is based on row/column similarity. It works as follows: a) it first divides the input table into blocks using the spanned cells as boundaries; b) it then compares how similar the last row/column in each block is to the previous ones using string, named-entity, and number similarity functions; c) then the right-most and/or bottom-most rows/columns that are similar to the last row/column are considered to contain data cells and the others are considered to contain meta-data cells. The evaluation was not performed on this task, but on their whole system.

Yoshida et al. [69] suggested using ontologies. Their proposal works as follows: a) for each cell in a table, it computes the ratio of times that its content is recorded in the ontology; b) these ratios are then used to feed the Expectation-Maximisation algorithm in order to learn a classifier that makes a few subtypes of listings apart; c) once the exact type of listing is clear, identifying meta-data cells is relatively easy and the rest of cells are assumed to be data cells. (Note that the authors assume that the input tables are data tables, which is the reason why this cannot be considered a discrimination proposal.) They evaluated their proposal on 175 tables that were randomly sampled from a repository with 35 232 tables.

Yang and Luk [68] applied some heuristics to differentiate rows with meta-data cells from rows with data cells. Their proposal works as follows: a) a row is considered to have meta-data cells if it has at most 50% the average number of cells per row, if it contains no structured cells, or if the visual features are different from the visual features of the others rows; b) then, it tries to detect if the input table is a listing or a matrix; c) once the table structure is identified, it is easy to identify the meta-data. (Note that the authors assume that the input tables are data tables, which is the reason why this cannot be considered a discrimination proposal.) The authors did not report on their experimental results regarding this task, but their whole system.

Kim and Lee [32] devised a proposal that first attempts to classify the input table. It works as follows: a) in the case of tables with one single row or column, the first cell is considered to be a meta-data cell and the rest are considered to be data cells; b) in the case of tables with two rows and two columns that

do not have any spanned cells, both the first row and column are considered to have meta-data cells and the bottom-right cell is considered to be a data cell; c) tables with two rows/columns and three or more columns/rows whose upper-left cell spans a whole row/column are discarded; otherwise if the first row/column has some spanned cells (but not all), then the first column/row is assumed to have meta-data cells and the others are assumed to have data cells; d) otherwise, the similarity of the cells is checked per rows and columns using the following functions: a lexical similarity function that focuses on the data types and the length of the contents, and a semantic similarity function that builds on some user-provided key words and patterns. The authors did not provide any experimental results regarding this task.

Jung and Kwon [30] proposed a heuristic-based technique to locate the meta-data within the tables. Their proposal works as follows: a) cells with a `th` tag are assumed to have meta-data; b) if the table can be partitioned into two blocks with the same background colour or font, then the top and/or the left blocks are assumed to contain meta-data; c) if the cells in a row or column have some user-defined contents or match some user-defined patterns, then they are also considered to contain meta-data; d) spanned cells that are embedded in `td` tags are also assumed to have meta-data as long as they are located on the top-left areas of the table; its adjacent cells are also considered to have meta-data; e) if the top-right cell is empty, then it is likely that the cells in the first row or column have meta-data; f) a probability is finally computed for every cell building on the previous heuristics and the cells whose probability exceeds a threshold are then considered to be meta-data cells whereas the others are assumed to be data cells. The evaluation was performed using 10-fold cross evaluation on 10 000 tables from their own repository plus the tables from Wang and Hu's [64] repository.

Cafarella et al. [3] devised a machine-learning proposal. It works as follows: a) a learning set is assembled with data tables in which the cells are classified as either meta-data or data cells; b) in cases in which a table does not have any meta-data cells, synthetic cells are created and the meta-data is fed from a separate database with similar tables; c) some structural and content features are computed for each cell; d) a classifier is machine-learned from the previous features; e) the results of the classifier are used to enrich the other database. The authors evaluated their proposal by means of 5-fold cross evaluation on a repository with 1 000 tables that were gathered from the Web.

Embley et al. [22] devised a heuristic-based proposal that searches for four critical cells that help delimit where the meta-data and the data cells are located. These cells are referred to as CC1, CC2, CC3, and CC4. CC1 and CC2 identify the top left-most region such that the cells on the right and below that region are mostly meta-data cells; CC3 and CC4 identify the bottom right-most region whose cells are mostly data cells. (Note that CC4 is identified in their segmentation task.) Their proposal works as follows: a) it sets CC1 to the top left-most cell and CC2 to the bottom left-most cell; b) it then iteratively shifts CC2 upwards or rightwards while searching for the minimum set of cells between CC1 and CC2 that result in headers that can identify the cells between

CC2 and CC4; c) then, CC3 is set to the first cell below to the right of CC2 that does not belong to an empty row or column; d) after that, footnotes are identified in cells whose contents start with a footnote-mark symbol; e) finally, it analyses some dependencies amongst the meta-data cells to find out the order in which they must be grouped. The evaluation was performed on a repository with 199 tables that was provided by Padmanabhan et al. [48].

Milošević et al. [44] restricted their attention to tables from the PubMed Central repository. Their focus is on identifying the meta-data cells, since the other cells are considered data cells by default. Their proposal works in three phases. In the first phase, it searches for `thead` tags; if they are found, then the inner `th` tags are assumed to encode the meta-data cells and their procedure finishes. Otherwise, the second phase is intended to find meta-data cells at the top rows as follows: a) they examine the syntactic similarity of cells on a per column basis; the cells at the top whose syntactic type is different from the cells below, if any, are considered meta-data cells as long as the cells in the same rows in adjacent columns are also considered meta-data cells; b) if a cell in the first row spans several columns, then it is assumed to have meta-data, as well as the cells in the rows below, until a non-spanned cell is found; c) the cells at the top that are between horizontal lines are considered meta-data if they are marked with a `thead` tag and they are not empty; in cases in which only one cell in a row has meta-data, the authors refer to it as a super row. The third phase is intended to find meta-data cells on the left columns as follows: a) the cells on the left-most column that are spanned are meta-data cells and so are the cells on the right until the first non-spanned cell is found; b) the first column below a super row is considered to have meta-data cells that are referred to as stubs. They used a repository with 3 573 tables from which 101 tables were randomly selected to evaluate their proposal.

### 3.5. Structural analysis

Penn et al. [49], Hurst [28], Wang and Hu [64], Kim and Lee [32], Jung and Kwon [30], Gatterbauer et al. [26], Okada and Miura [47], Cafarella et al. [3], Crestan and Pantel [16], Lautert et al. [34], Son and Park [59], Eberius et al. [18], Embley et al. [22], Nishida et al. [46], and Liao et al. [38] did not report on any proposals to implement the structural analysis task. Chen et al. [7] presented a naive proposal that works on tables that provide data about a single entity, so all of the data cells form a single tuple; regarding the meta-data cells, they group them into headers horizontally or vertically after expanding spanned cells. Yoshida et al. [69] presented a naive proposal that classifies tables in a number of categories, which makes identifying the tuples quite a trivial task. Elmeleegy et al. [20] also assumed that the tuples within tables that are encoded as lists are always laid out row-wise. Ling et al. [39] and Braunschweig et al. [1] assumed that tuples are displayed row-wise or column-wise depending on the number of meta-data or data cells found in the first few rows or columns. Chu et al. [9] also presented a naive approach that assumes that the tuples within tables that are encoded as lists are always laid out row-wise. Wu et al. [66] presented an

additional naive approach since they just identify tuples in horizontal listings. The other proposals provide more sophisticated approaches.

Lerman et al. [36] used a couple of algorithms to detect row-wise tuples. Their proposal works as follows: a) first, it uses DataPro to find the patterns that describe the data in each column; b) such patterns can be interpreted as tags that allow to transform a table into a sequence of symbols; c) then, a version of ALERGIA is used to infer a finite automaton from those sequences; d) the automaton is then transformed into a regular expression; e) finally, it identifies repeating sub-patterns that correspond to the tuples in the original table. No experimentation was performed regarding this task.

Cohen et al. [13] presented a proposal that relies on four so-called builders, namely: a builder focuses on meta-data cells that cut in on the table, one that focuses on columns of headers, another that focuses on rows of headers, and an additional one that takes the tag paths into account. The builders are fed into a FOIL-based system in order to learn a classification rule that allows to identify both horizontal and vertical tuples. No experimental results were reported regarding this specific task, but their whole system.

Yang and Luk [68] presented a proposal that specialises in numerical tables. It works as follows: a) first, it removes the headers of the input table; b) then, it checks whether the tuples seem to be one-dimensional or two-dimensional using some heuristics; c) the type of cells is analysed using pre-defined patterns in order to label numeric data cells; d) given the types of cells and the dimensionality of the tuples, their proposal tries to match a number of pre-defined patterns that help identify the tuples. The evaluation was performed on 169 one-dimensional and 50 two-dimensional tables.

Lerman et al. [37] devised two proposals to identify tuples, namely: a constraint solving technique and a probabilistic technique. The former works as follows: a) it models the cells in the tables using Boolean variables; b) it then adds constraints to ensure that each cell belongs to a single tuple, only contiguous cells can be assigned to the same tuple, and two cells cannot be in the same position in the same table; c) then a constraint solver is used to find a solution to the constraints. The latter works as follows: a) it uses a set of observable variables that model the types of tokens in the data cells, and a set of hidden variables, which provide the tuple number or the column number to which every cell belongs; b) a probabilistic model is then learnt by assuming a number of dependencies between token types, cells, columns, neighbour columns, format, or tuple numbers; c) finally, the contents of the hidden variables are inferred building on the probabilistic model. Their evaluation was performed on the tables from their own repository, which were gathered from 12 web sites on book sellers, property taxes, white reports, and prisons.

Fumarola et al. [24] presented a proposal that was described very shallowly. It seems to work on so-called candidate lists, which are sets of cells that correspond to different columns and form a single tuple; each candidate list is a sub-tree of the DOM tree and they all are required to satisfy some structural similarity constraints, including a minimum size in terms of nodes. The evaluation was performed on 224 tables from Gatterbauer et al.'s [26] repository.

Milošević et al. [44] identify the tuples according to how the meta-data cells are arranged within a table. If meta-data cells are at the top-most rows and on the left-most columns, then, the table is a matrix with a single tuple that consists of the whole set of data cells; if there are meta-data cells at the top, but not on the left, then the table is a listing in which each row is a tuple; if there are not any meta-data cells, every single data cell corresponds to a tuple. (Note that this proposal cannot be considered a discrimination proposal since it assumes that the input tables are data tables; recall that their focus was on tables in PubMed publications, which are tables with scientific data.) They used a repository with 3 573 tables from which 101 tables were selected to evaluate their proposal.

### 3.6. Interpretation

Lerman et al. [36], Penn et al. [49], Yoshida et al. [69], Cohen et al. [13], Hurst [28], Wang and Hu [64], Lerman et al. [37], Kim and Lee [32], Jung and Kwon [30], Gatterbauer et al. [26], Okada and Miura [47], Crestan and Pantel [16], Elmeleegy et al. [20], Fumarola et al. [24], Lautert et al. [34], Ling et al. [39], Son and Park [59], Braunschweig et al. [1], Chu et al. [9], Eberius et al. [18], Nishida et al. [46], and Liao et al. [38] did not report on this task.

Most of the other authors reported on naive solutions. Chen et al.'s [7] proposal works on tables with a single tuple that can spread across several blocks, each of which has its own headers; for each component of the tuple, it creates field descriptors by joining the meta-data cells in the corresponding rows and/or columns. Embley et al.'s [22] proposal is similar, but they focused on tables with a single block. Milošević et al. [44] reported on a naive approach, too: in matrices or listings, they create descriptors for each component from the meta-data in the corresponding column and/or row; in enumerations, they use the caption of the table as a descriptor for every component in the tuples. Yang and Luk [68] proposed a similar procedure, but it takes multiple header rows or columns into account, in which case the cells are simply merged to create field descriptors, as well as cells that contain both meta-data and data, in which case the meta-data are transformed into simple descriptors.

The proposal by Cafarella et al. [3] goes a step forward in cases in which a table does not provide any meta-data cells. In such cases, they collect the data on a per-column basis and attempt to find the most similar data in the ACSDB database, which is a resource that has many data with correct descriptors. The authors did not report on the evaluation of this task. Wu et al. [66] went also a bit further since they used several ad-hoc interpretation methods depending on the structure of the table identified in the discrimination task. They only reported on a method to extract information from horizontal listings with headers using some heuristics that are related to how the `th` and the `td` tags encode a subject-predicate-object relation. They conducted their experimentation on a repository with 100 horizontal listings from Wikipedia. The authors evaluated their proposal on a repository with 5 000 tables that were randomly selected from the Wikipedia.

## 4. Comparison of proposals

In this section, we compare the proposals that we have summarised in the previous section by means of a comparison framework with both general and task-specific characteristics.

The general characteristics are the following: a) *Foundation*: it is a hint on the technique behind each proposal. b) *Tables required*: it is the minimum number of tables required for a proposal to work; the less tables required, the better. c) *Effectiveness*: it is the extent to which a proposal succeeds in implementing a task correctly according to an effectiveness measure; the higher the effectiveness, the better. d) *Efficiency*: it is the amount of computing power that a proposal requires to implement a task; the more efficient (i.e., the less computing power is required), the better. e) *Resources*: it refers to the resources that a user must provide so that a proposal can work properly; the less resources, the better. f) *Features*: it refers to the features onto which the input data must be projected in order to machine learn a predictor or to make a decision according to a heuristic. Features can be either structural, which are related to the HTML or the DOM representation of the input documents, visual, which are related to how they are displayed, or content features, which are related to the contents of the cells. g) *Parameters*: it refers to the settings that must be tuned so that a proposal works well, which can be either pre-defined, learnable, or user-defined parameters. Pre-defined parameters have a value that the authors of a proposal have found generally appropriate; they are preferable to learnable parameters, whose values must be experimentally learnt by the user; in turn, they are preferable to user-defined parameters, which must be set by the user using his or her intuition; the less parameters, the better.

Note that it is easy to make decisions building on the general features that we presented above since we have characterised their preferred contents; the same applies to the task-specific features that we describe in the following subsections. The only exceptions are the foundation characteristic and the features characteristic. The reason is that it is not generally clear whether a heuristic-based approach is preferable to a machine-learning approach or vice versa, or whether structural, visual, or content features are preferable to each other. Note, too, that effectiveness and efficiency are decision-making characteristics, but the figures provided by an author are not generally comparable to the figures provided by a different author because they evaluated their proposals using different approaches, learning sets, evaluation sets, and machinery.

### 4.1. Location

Table 2 summarises our comparison regarding location proposals. The task-specific characteristics are the following: a) *Body encodings*: it refers to how the tables that a proposal can locate must be encoded; the more kinds of encodings are identified, the better. b) *Context-data encodings*: it refers to how context-data cells are encoded; the more kinds of encodings are identified, the better.

Regarding the general characteristics, it is surprising that all of the location proposals are based on heuristics; there is no record in the literature of a single

Table 2: Comparison of location proposals.

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Body encodings	Context data encodings
						Structural	Visual	Content	Pre-defined	Learnable	User-defined		
Chen et al. 2000	Heuristics	1	N/A	N/A								Tabular tags	Tag caption
Lerman et al. 2001	Heuristics	2+	N/A	N/A								Listing tags	
Penn et al. 2001	Heuristics	1	N/A	N/A								Leaf tabular tags	
Cohen et al. 2002	Heuristics	1	N/A	N/A								Tabular tags	
Hurst 2002	Heuristics	1	N/A	N/A								Tabular tags	
Wang and Hu 2002	Heuristics	1	N/A	N/A								Leaf tabular tags	
Yang and Luk 2002	Heuristics	1	N/A	N/A								Tabular tags	
Lerman et al. 2004	Heuristics	2+	N/A	N/A								Listing tags	
Kim and Lee 2005	Heuristics	1	N/A	N/A								Tabular tags	
Jung and Kwon 2006	Heuristics	1	N/A	N/A								Tabular tags	
Gatterbauer et al. 2007	Heuristics	1	N/A	N/A		Presence of tag table; cell adjacency.	CSS attributes color, bgcolor, font-size, font-style, font-weight, font-family, and text-align; attribute href.		Maximum distance to adjacent cells			Any	
Okada and Miura 2007	Heuristics	1	N/A	N/A								Tabular tags	
Cafarella et al. 2008	Heuristics	1	N/A	N/A								Tabular tags	
Crestan and Pantel 2011	Heuristics	1	N/A	N/A						Minimum number of samples per node		Leaf tabular tags	
Fumarola et al. 2011	Heuristics	1	N/A	N/A								Any	
Ling et al. 2013	Heuristics	2+	N/A	N/A								Tabular tags	Tag title; surrounding text.
Son and Park 2013	Heuristics	1	N/A	N/A								Tabular tags	
Chu et al. 2015	Heuristics	1	N/A	N/A								Listing tags	
Eberius et al. 2015	Heuristics	1	N/A	N/A								Tabular tags	
Milošević et al. 2016	Heuristics	1	N/A	N/A								Tabular tags	
Wu et al. 2016	Heuristics	1	N/A	N/A								Tabular tags	
Liao et al. 2018	Heuristics	1	N/A	N/A								Tabular tags	

proposal that has tried a machine-learning approach. Most proposals can work on a single table, but the ones by Lerman et al. [36, 37] and Ling et al. [39] require at least a pair of tables to perform table alignment. None of the proposals was presented in isolation, but as a component of a larger system, which is the reason why no author reported on effectiveness or efficiency. Realise that only the proposal by Gatterbauer et al. [26] projects the input documents onto structural and visual features in order to apply their heuristics; note, too, that it is the only one that requires a pre-defined parameter. The proposal by Crestan and Pantel [16] is the only that requires the user to set a learnable parameter.

Regarding the task-specific characteristics, most of the proposals locate tables that are encoded using tabular tags, a few focus on tables that are encoded using listing tags, and only Gatterbauer et al.’s [26] and Fumarola et al.’s [24] proposals are independent from the tags used since they analyse how the input documents are displayed. Note, too, that the vast majority of proposals focus on locating the tables themselves, not their context data. Chen et al. [7] and Ling et al. [39] are the exceptions: the former presents a simple approach that searches for caption tags and the latter presents a more sophisticated approach that analyses the title tags and the text that surrounds the tables.

#### 4.2. Segmentation

Table 3 summarises our comparison regarding segmentation proposals. The task-specific characteristics are the following: a) *Spanned cells*: it describes if a proposal is able to identify cells that span multiple columns and/or rows; a proposal that can identify spanned cells is better than a proposal that cannot. b) *Multi-part cells*: it describes if a proposal is able to identify cells that provide partial contents and must be merged; a proposal that can identify multi-part cells is better than a proposal that cannot. c) *Context data*: it describes if a proposal can identify context data or not; a proposal that can identify context data is better than a proposal that cannot.

Regarding the general characteristics, it is easy to realise that only the proposals by Lerman et al. [36, 37] have tried machine-learning approaches; the others rely on heuristics that their authors have proven to work well in practice. Furthermore, most of them can work on as few as one input table, but the ones by Lerman et al. [36, 37] and Ling et al. [39]. Unfortunately, roughly 72% of the authors did not report on the effectiveness of their proposals; the others reported on precision, recall, and/or the  $F_1$  score. Only Elmeleegy et al. [20] and Chu et al. [9] reported on the efficiency of their approaches; their figures reveal that the algorithms behind the scenes might not be scalable enough. Regarding the resources required, only the proposals by Elmeleegy et al. [20] and Ling et al. [39] require the user to provide a few, but they do not seem to be difficult to find. Only the proposals by Lerman et al. [36, 37] and Gatterbauer et al. [26] require to project the input tables onto some simple features. Regarding the parameters, only the proposals by Elmeleegy et al. [20], Ling et al. [39], and Chu et al. [9] need the users to set a few.

Regarding the task-specific characteristics, it is surprising that many proposals do not make an attempt to analyse spanned cells and that none of them



Table 3: Comparison of segmentation proposals.

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Spanned cells	Multi-part cells	Context data
						Structural	Visual	Content	Pre-defined	Learnable	User-defined			
Chen et al. 2000	Heuristics	1	N/A	N/A								Yes	No	Yes
Lerman et al. 2001	Clustering; probabilistic model.	2+	N/A	N/A				DataPro patterns				No	No	No
Cohen et al. 2002	Heuristics	1	N/A	N/A								Yes	No	No
Yang and Luk 2002	Heuristics	1	N/A	N/A								Yes	No	No
Lerman et al. 2004	Clustering	2+	N/A	N/A				Token types				No	No	No
Jung and Kwon 2006	Heuristics	1	N/A	N/A								Yes	No	No
Gatterbauer et al. 2007	Heuristics	1	N/A	N/A				Bounding boxes; padding; borders; margins.				Yes	No	No
Elmeleegy et al. 2011	Heuristics	1	F <sub>1</sub> =0.92	Average 8 seconds per list	Language model; corpus with 154M tables.				Min TC support	Best-first search nodes	Maximum number of repetitions	No	No	No
Ling et al. 2013	Heuristics	2+	P=0.62 R=0.70 F <sub>1</sub> =0.66	N/A	Wikipedia corpus						Same heuristic score; empty segment score.	No	No	Yes
Chu et al. 2015	Heuristics	1	<b>Wiki</b> P=0.89 R=0.93 F <sub>1</sub> =0.90 <b>Other</b> P=0.89 R=0.93 F <sub>1</sub> =0.90	Average 2 seconds per row					Refinement required; syntactic distance weight; semantic distance weight.		Tokenisation delimiters	Yes	No	No
Embley et al. 2016	Heuristics	1	N/A	N/A					Minimum number of empty cells in context data rows			Yes	No	Yes

can identify multi-part cells, both of which are very common in practice. It is also surprising that only the proposals by Chen et al. [7], Ling et al. [39], and Embley et al. [21] can identify context data, which are also very common in practice; unfortunately, the proposal by Chen et al. [7] cannot be considered a general solution to the problem since it is very naive.

#### 4.3. Discrimination

Tables 4–7 summarise our comparison regarding discrimination proposals. The only task-specific characteristic is *Types of data tables*, which refers to the kinds of data tables that a proposal can discriminate; the more types can be discriminated, the better.

Regarding the general characteristics, it is easy to realise that 64% of the proposals use a machine-learning approach and the rest use heuristic-based approaches. The former require at least two tables to learn a predictor that implements the discrimination task, whereas the latter can generally work on a single table. Except for Wu et al.’s [65], the other authors report on effectiveness measures that are specific to this task; most of the authors selected precision, recall, and the  $F_1$  score as effectiveness measures; the exceptions are Cohen et al. [13], Lautert et al. [34], and Nishida et al. [46], who report on the  $F_1$  score only, Okada and Miura [47], who reported on accuracy, and Fumarola et al. [24], who reported on recall only. Apparently, the effectiveness of the machine-learning proposals is higher than the effectiveness of the heuristic-based proposals; however, due to the differences in the evaluation processes, this conclusion is not sound. Unfortunately, only Son and Park [59] and Eberius et al. [18] reported on the efficiency of their proposals, which does not seem to be very good according to their figures; Wu et al. [65] did not report on the efficiency of their proposal but they mentioned that it relies on a linear clustering algorithm. The only proposals that require resources are the ones by Eberius et al. [18] and Nishida et al. [46]; fortunately, they do not seem to be a major obstacle since they consist in a corpus that was gathered from the Wikipedia. The ones that rely on machine learning project the input data onto a space of structural, visual, and/or content features that seem simple to compute. Regarding their parameters, most of them have pre-defined parameters for which the authors recommend some values that are expected to work generally well; none of the proposals require any learnable parameters, but a few require user-defined parameters.

Regarding the task-specific characteristics, the only proposals that can subclassify data tables are the following ones: Crestan and Pantel [16] distinguishes amongst listings, forms, matrices, and enumerations; Lautert et al. [34], Eberius et al. [18], and Nishida et al. [46] distinguish amongst listings, forms, and matrices; and Liao et al. [38] distinguishes between complete and incomplete tables (which are encoded as independent tables, but must be stitched together so that they can be properly interpreted).

#### 4.4. Functional analysis

Table 8 summarises our comparison regarding functional analysis proposals.

Table 4: Comparison of discrimination proposals (Part 1).

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Types of data tables
						Structural	Visual	Content	Pre-defined	Learnable	User-defined	
Chen et al. 2000	Heuristics	1	P=0.92 R=0.80 F <sub>1</sub> =0.86	N/A					Thresholds to the number of links, forms, figures, and neighbour similar cells; string similarity threshold; named entity similarity threshold; number category similarity threshold.			
Penn et al. 2001	Heuristics	1	P=0.86 R=0.89 F <sub>1</sub> =0.88	N/A						Average word length threshold.		
Cohen et al. 2002	Winnow	2+	F = 0.79-1.00	N/A	Number of rows; number of columns.			Ratio of columns with alphabetic content; ratio of simple cells.		Winnow threshold.		
Hurst 2002	Naive Bayes; Winnow.	2+	<b>Naive Bayes</b> P=0.95 R=0.94 <b>Winnow</b> P=1.00 R=0.92 F = 0.96	N/A	Number of tr tags; maximum number of th/td tags; bag-of-tags beneath the table tag.	Border; bag-of-attributes of the tags inside a table tag; number of rows or columns according to a geometric model		String content ratio; singular cell ratio.				
Wang and Hu 2002	Custom decision tree learner	2+	P=0.97 R=0.94 F <sub>1</sub> =0.95	N/A	Average and standard deviation of the number of cells, cells per row and cells per column; average cumulative length consistency.			Histogram and consistency of each value type; word group frequency features.	Naive Bayes parameters; maximum impurity reduction; maximum depth of the tree; minimum number of samples used.			
Yang and Luk 2002	Heuristics	1	P=0.94 R=1.00 F <sub>1</sub> =0.97	N/A						Attribute-indicating entity patterns Value-indicating entity patterns.		
Kim and Lee 2005	Heuristics	1	P=0.97 R=0.99 F <sub>1</sub> =0.98	N/A					Table homogeneity threshold; syntactic homogeneity threshold.	Threshold to the number of link/image/empty cells; homogeneity patterns.		

Table 5: Comparison of discrimination proposals (Part 2).

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Types of data tables
						Structural	Visual	Content	Pre-defined	Learnable	User-defined	
Jung and Kwon 2006	C4.5	2+	P=0.94 R=0.96 F <sub>1</sub> =0.95	N/A		Presence of a caption, th, or table tag; tag type; ratio of spanned cells; likelihood of headers; index based on the number of columns and rows; tag consistency.	Border	Presence of rows and columns with numerical data; ratio of empty, string-only, symbol-only, and digit-only cells; ratio of cells with images or links; number of sentences longer than 40 characters; average standard deviation of the cell length; consistency of content type.				
Gatterbauer et al. 2007	Heuristics	1	P=0.68 R=0.81 F <sub>1</sub> =0.74	N/A					Maximum cell coverage; maximum word length.			
Okada and Miura 2007	ID3	2+	Acc=0.91	N/A		Presence of nested tables, images, spanned cells, or th tag; number of HTML tags and rows.	Presence of visible border lines or caption tag; specification of width or height.					
Cafarella et al. 2008	Statistic model	2+	P=0.41 R=0.81 F <sub>1</sub> =0.54	N/A		Dimensions		Ratio of rows with mostly void cells and non-string data; average cell length and standard deviation; number of columns with basic types.				
Crestan and Pantel 2011	Gradient Boosted Decision Tree	2+	P=0.80 R=0.76 F <sub>1</sub> =0.78	N/A		Maximum number of cells per row and column; ratio of cells with th, a, img, input, select, br, or formatting tags; ratio of colspans and rowspans.		Maximum cell length; average and variance of cell lengths; ratio of distinct tags, strings, cells with trailing colons, cells with digits, numerical cells, and non-empty cells.	Number of trees; minimum samples per node, best-first search nodes at classification.			Listings; forms; matrices; enumerations.
Fumarola et al. 2011	Heuristics	2+	R=0.79	N/A						Value similarity threshold; structure similarity threshold.		

Table 6: Comparison of discrimination proposals (Part 3).

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Types of data tables
						Structural	Visual	Content	Pre-defined	Learnable	User-defined	
Lautert et al. 2013	Neural networks	2+	$F_1=0.22-0.95$	N/A		Maximum number of cells per row and column; ratio of cells with th, a, img, input, select, br, or formatting tags; colspan vs. rowspan ratio; type of table.	Position of inner tables	Maximum cell length; average and variance length; ratio of distinct tags, trailing colon cells, cells with digits, numerical cells, non-empty cells; ratio of cells with commas and brackets; ratios of cells with unorderd and ordered lists.	Number of hidden layers.			Listings; forms; matrices.
Son and Park 2013	SVM	2+	$P=0.98$ $R=0.98$ $F_1=0.98$	<b>Learning</b> 20h <b>Evaluation</b> 2h (1100 tables)		Average and standard deviation of the number of cells, cells per row and cells per column; average cumulative length consistency.		Histogram and consistency of each value type; word group frequency features.	Parse tree kernel decay factor; structure kernel mixing ratio; composite kernel mixing ratio.			
Eberius et al. 2015	Random Forest	2+	<b>One classifier</b> $P=0.84$ $R=0.82$ $F_1=0.82$ <b>Two classifiers</b> $P=0.83$ $R=0.83$ $F_1=0.83$	10,238h (millions of tables)	Wikipedia	Maximum, average, and standard deviation of row, column, and cell lengths; ratio of images, forms, hyperlinks, and other cells; presence of th tags; local cell length average and variance; local ratio of span, th, a, img, input, select, font, br, ul, or ol tags.		Ratio of alphabetic, digit, and empty cells; local ratio of colons, commas, brackets, numeric characters, and non-empty cells.	Local regions			Listings; forms; matrices.

Table 7: Comparison of discrimination proposals (Part 4).

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Types of data tables
						Structural	Visual	Content	Pre-defined	Learnable	User-defined	
Wu et al. 2016	Clustering	2+	N/A	N/A		Similarity to other tables; local cluster density.					Cut-off distance	
Nishida et al. 2017	Deep neural network	2+	<b>Single</b> F =0.88 <b>Ensemble</b> F <sub>1</sub> =0.91	N/A	Wikipedia				Token embedding size; LSTM output size; number of residual units; number of fully connected layers; number of weighted layers; CNN filters; SGD momentum; SGD minibatch size.			Listings; forms; matrices.
Liao et al. 2018	Heuristics	1	P=0.94 R=0.93 F <sub>1</sub> =0.94	N/A					Empty cell threshold		Layout tag threshold; consistency threshold.	Complete; incomplete.



The task-specific characteristics are the following: a) *Context-data cells*: it describes if a proposal is able to make context-data cells apart from the others; a proposal that can identify context-data cells is better than a proposal that cannot. *Decorators*: it refers to the ability of a proposal to identify decorator cells; a proposal that can find decorators is better than a proposal that cannot.

Regarding the general characteristics, 83% of the proposals rely on heuristics and rest rely on machine-learning approaches. Most of them can work on as few as a single table, with the exception of the proposals by Cafarella et al. [3], Chen et al. [7], and Chen et al. [7], we need to compare at least two tables. Many of the authors report on the effectiveness of their proposals; realise that most of the measures are below 0.90, which means that there is enough room for improvement regarding this task. Unfortunately, only Embley et al. [21] reported on the efficiency of their proposal, which seems scalable enough. Regarding the resources required, Yoshida et al.’s [69] and Gatterbauer et al.’s [26] proposals require domain-specific ontologies, whereas Cafarella et al.’s [4] requires a publicly-available database. The proposal by Cafarella et al. [4] is the only that projects the input data onto a space of simple structural and content features. The proposal by Yoshida et al. [69] requires a pre-defined parameter that is auto-adjusted, and the proposals by Yang and Luk [68], Kim and Lee [32], and Milošević et al. [44] require another pre-defined parameter for which the authors provide a default value; the only proposals that require user-defined parameters are the ones by Chen et al. [7], Kim and Lee [32], Jung and Kwon [30], and Ling et al. [39].

Regarding the task-specific characteristics, note that only the proposal by Milošević et al. [44] can identify some decorator cells and context-data cells. This is a bit surprising since, according to our experience, these kinds of cells are very common in practice.

#### 4.5. Structural analysis

Table 9 summarises our comparison regarding structural analysis proposals. The task-specific characteristics are the following: a) *Header structure*: it describes the kinds of headers that a proposal can identify according to their structure, namely: **none**, which means that it can analyse tables without headers, **simple**, which means that it can analyse simple headers that consists of one meta-data cell only, and **complex**, which means that it can identify complex headers that consists of multiple meta-data cells; the more header structures a proposal can identify, the better. b) *Header layout*: it describes the kinds of headers that a proposal can identify according to how they are laid out, namely: **none**, which means that it can identify that a table does not have any headers, **single**, which means that it can identify headers in the first rows and/or columns of a table, **horizontally repeated**, which means that it can identify that the headers are repeated every some rows, **vertically repeated**, which means that it can identify that the headers are repeated every some columns, and **split**, which means that it can identify series of headers that are split across several non-adjacent rows or columns; the more header layouts a proposal can identify, the better. c) *Tuple dimensionality*: it describes the dimensionality of the tuples



Table 9: Comparison of structural analysis proposals.

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Header structure	Header layout	Tuple dimensionality	Tuple multiplicity	Tuple orientation	Separators
						Structural	Visual	Content	Pre-defined	Learnable	User-defined						
Chen 2000	Heuristics	1	N/A	N/A				String similarity; entity similarity; number category similarity.			Thresholds for similarity	Simple; complex.	Single	1, 2	1, *	None; horizontal; vertical.	No
Lerman et al. 2001	DataPro Grammar induction; optimisation.	2+	N/A	N/A								None	None	1	*	Horizontal	No
Yoshida et al. 2001	Heuristics	2+	N/A	N/A								Simple	None; single; split.	1	*	Horizontal; vertical.	No
Cohen et al. 2001	Inductive logic programming	1	N/A	N/A								Simple	Single	1	*	Horizontal; vertical.	No
Yang and Luk 2002	Heuristics	1	<b>1 dim.</b> P=0.92 R=0.96 F <sub>1</sub> =0.94 <b>2 dim.</b> P=0.90 R=0.98 F <sub>1</sub> =0.94	N/A				Average cell count threshold			Cell type patterns	None	None	1, 2	*	Horizontal; vertical.	No
Lerman et al. 2004	Constraint Satisfaction; Hidden Markov Models.	2+	N/A	N/A								None	None	1	*	Horizontal	No
Elmeleegy et al. 2011	Heuristics	1	P=0.64 R=0.63 F <sub>1</sub> =0.63	N/A								Simple	Single	1	*	Horizontal	No
Fumarola et al. 2011	Heuristics	2	N/A	N/A							Minimum nodes threshold	None	None	0	*	None	No
Ling et al. 2013	Heuristics	1	N/A	N/A								Simple	Single	1	*	Horizontal	No
Braunschweig et al. 2015	Heuristics	1	N/A	N/A								Simple	Single	1	*	Horizontal	No
Chu et al. 2015	Heuristics	1	N/A	N/A								None	None	1	*	Horizontal	No
Milošević et al. 2016	Heuristics	1	P = 0.92 R = 0.97 F <sub>1</sub> = 0.94	N/A								None; simple; complex.	None; single; split.	0, 1, 2	*	None; horizontal.	No
Wu et al. 2016	Heuristics	1	N/A	N/A								Simple	Single	1	*	Horizontal	No

that a proposal can identify, namely: 0 if it can identify the tuples in an enumeration, 1 if it can identify the tuples in a listing or a form, and 2 if it can identify the tuple in a matrix; the more tuple dimensionalities a proposal can identify, the better. d) *Tuple multiplicity*: it describes the number of tuples that a table is intended to show, namely: 1 in the case of forms and matrices, and \* in the case of listings and enumerations; the more tuple multiplicities a proposal can identify, the better. e) *Tuple orientation*: it describes the orientations that it can identify, namely: none in the case of matrices and enumerations, horizontal or vertical in the case of listings and forms; the more tuple orientations a proposal can identify, the better. f) *Separators*: it describes whether a proposal can identify separator rows and/or columns; a proposal that can identify separators is better than a proposal that cannot.

Regarding the general characteristics, many proposals rely on heuristic-based approaches; the exceptions are the proposals by Lerman et al. [36, 37], which leverage some grammar induction techniques, and Cohen et al.’s [13] proposal, which leverages inductive logic programming. Most of the proposals require as few as one input table; the exceptions are the proposals by Lerman et al. [36, 37], Yoshida et al. [69], and Fumarola et al. [24], which require two tables for comparison purposes. Unfortunately, only Yang and Luk [68], Elmeleegy et al. [20], and Milošević et al. [44] reported on the effectiveness of their proposals, and none of the authors reported on their efficiency. Note that none of the proposals require to project the input data onto a space of features, but the one by Chen et al. [7]. Note, too, that Chen et al. [7], Yang and Luk’s [68], and Fumarola et al.’s [24] proposals are the only that have parameters.

Regarding the task-specific characteristics, it is surprising that most of the proposals assume that the tables do not have any headers or they are simple, except for Milošević et al.’s [44] proposal; it is also surprising that the only proposal that can identify single and split headers is the one by Yoshida et al. [69]. Regarding the tuple dimensionality, only the proposals by Yang and Luk [68] and Milošević et al. [44] can make uni-dimensional tuples apart from two-dimensional tuples; Milošević et al.’s [44] can also deal with zero-dimensional tuples; the proposal by Fumarola et al. [24] implicitly assumes that the tuples in a table are zero-dimensional and does not make an attempt to analyse the structure of the corresponding cells; the other proposals implicitly assume that the tuples are uni-dimensional. Regarding the tuple multiplicity, it is interesting to see that all of the proposals assume that tables may display more than one tuple; simply put, they cannot make listings apart from forms. Regarding the tuple orientation, most proposals implicitly assume that the tuples are oriented horizontally; the only exceptions are the proposals by Yoshida et al. [69], Cohen et al. [13], and Yang and Luk [68], which can make horizontal tuples apart from vertical tuples. It is surprising that none of the proposals that we have surveyed can identify separators, even though they are very common in practice.

#### 4.6. Interpretation

Table 10 summarises our comparison regarding interpretation proposals. The task-specific characteristics are the following: a) *Descriptors*: it reports

Table 10: Comparison of interpretation proposals.

Reference	Foundation	Tables required	Effectiveness	Efficiency	Resources	Features			Parameters			Descriptors	Empty values	Value structure
						Structural	Visual	Content	Pre-defined	Learnable	User-defined			
Chen et al. 2000	Heuristics	1	N/A	N/A								Simple; field.	No	No
Yang and Luk 2002	Heuristics	1	N/A	N/A								Simple; field.	No	Yes
Cafarella et al. 2008	Reference matching	1	N/A	N/A	ACSDB database							Simple	No	No
Embley et al. 2016	Heuristics	1	N/A	N/A								Simple	No	No
Milošević et al. 2016	Heuristics	1	N/A	N/A								Simple; field.	No	No
Wu et al. 2016	Heuristics	1	P=0.66	N/A								Simple	No	No

on the kind of descriptors that a proposal can assign to the data in a table; the more kinds of descriptors a proposal can generate, the better. b) *Empty contents*: it refers to the ability of a proposal to make a difference between empty cells whose contents are factorised and cells that are actually empty; a proposal that can make a difference between factorised and void cells is better than another proposal that cannot. c) *Content structure*: it refers to the ability of a proposal to make a difference between atomic cells and structured cells; a proposal that can make a difference between atomic cells and structured cells is better than a proposal that cannot.

Regarding the general characteristics, most proposals rely on heuristics that have proven to work well in practice; the only exception is the proposal by Cafarella et al. [4], which uses a reference matching approach. Wu et al. [65] were the only authors who reported on effectiveness, but they measured precision only; unfortunately, none of the proposals report on efficiency. Cafarella et al.'s [4] proposal is the only one that requires a publicly-available resource. None of the proposals project the input data onto a feature space and none of them require any parameters to be set.

Regarding the task-specific characteristics, all of the proposals can generate simple descriptors; only the proposals by Chen et al. [7], Yang and Luk [68], and Milošević et al. [44] can generate field descriptors. Unfortunately, none of the proposals can make a difference between factorised cells and void cells. Regarding making a difference amongst atomic and structured cells, it seems that only the proposal by Yang and Luk [68] can deal with this problem.

## 5. Conclusions

This article summarises and compares many proposals that have been published between 2000 and 2018 regarding extracting data from tables that are encoded using HTML. The problem is not trivial insofar many tables are encoded using a subset of table-related tags that help locate and segment them, but do not provide a clue on the function of the cells or their structure; many others are encoded using listing tags, block tags, or other tags that look like a table when they are displayed, which hampers locating and segmenting them.

Our analysis makes it clear that none of the proposals that we have listed provide a complete solution to the data-extraction problem. Most of them address only some of the tasks involved and they differ regarding the problems that they address within each task. Regarding the location task, most proposals focus on tables that are encoded using table-related tags, there are a couple that focus on listing tags, and also a couple that are independent from the tags used; what seems an actual challenge is to identify context data, since the few proposals that take this problem into account are very naive. Regarding the segmentation task, it is surprising that no proposal can identify multi-part cells and that most of them do not attempt to segment the context data. The discrimination task is the one that has been paid more attention, but not many proposals attempt to go further than making non-data tables apart from data tables; recent proposals attempt to classify data tables in more categories since

this definitely helps interpret them. Regarding the functional analysis task, it is surprising that almost none of the proposals pay attention to identifying context-data cells or decorators cells. Regarding the structural analysis task, the problems that have got none or very little attention are identifying split headers and zero- and two-dimensional tuples. Regarding the interpretation tasks, creating artificial descriptors in cases in which not enough meta-data are available, analysing whether an empty value is actually empty or factorised, and analysing the structure of the contents of a cell are problems that have not been addressed sufficiently. Addressing these problems would help expand the kinds of tables from which data can be extracted.

Last, but clearly not least, the evaluation of the proposals is also a very relevant problem. We have found that many authors have used Wang and Hu’s [64] repository in addition to their own repositories; unfortunately, the subsets of tables selected were different and their sizes range from as many as 342 795 tables to a hundred tables or less. Definitely, recent repositories like DWDTTC [18] or WDC [35] will help. We have also found many authors who used  $k$ -fold cross evaluation, but there is not a general consensus; there is not even a consensus regarding the value of  $k$  in the cases in which this procedure was used. As a conclusion, the experimental results reported are not comparable to each other. Neither is it common to find figures regarding efficiency, which makes it difficult to realise if a proposal might work well in a production scenario. Jiménez et al. [29] set a foundation regarding how to evaluate information extraction proposals in general, but they did not focus on the tasks involved in extracting information from tables that are encoded using HTML.

Summing up: extracting data from tables that are encoded in HTML is an active research field in which we expect new results to be published in the near future. We hope that this article helps researchers sift through the state-of-the-art proposals in this field.

## Acknowledgments

The work by Juan C. Roldán, Patricia Jiménez, and Rafael Corchuelo was supported by the Spanish R&D programme with grants TIN2013-40848-R and TIN2016-75394-R. The work by Juan C. Roldán was also supported by the Fulbright programme.

## References

- [1] K. Braunschweig, M. Thiele, and W. Lehner. From web tables to concepts: a semantic normalization approach. In *ER*, pages 247–260, 2015. doi: 10.1007/978-3-319-25264-3\_18.
- [2] A. L. Buchsbaum, D. F. Caldwell, K. W. Church, G. S. Fowler, and S. Muthukrishnan. Engineering the compression of massive tables: an experimental approach. In *SODA*, pages 175–184, 2000. URL <http://dl.acm.org/citation.cfm?id=338219.338249>.

- [3] M. J. Cafarella, A. Y. Halevy, D. Z. Wang, E. Wu, and Y. Zhang. WebTables: exploring the power of tables on the Web. *PVLDB*, 1(1):538–549, 2008. URL <http://www.vldb.org/pvldb/1/1453916.pdf>.
- [4] M. J. Cafarella, A. Y. Halevy, Y. Zhang, D. Z. Wang, and E. Wu. Uncovering the relational Web. In *WebDB*, 2008. URL <http://webdb2008.como.polimi.it/images/stories/WebDB2008/paper30.pdf>.
- [5] M. J. Cafarella, A. Y. Halevy, H. Lee, J. Madhavan, C. Yu, D. Z. Wang, and E. Wu. Ten years of web tables. *PVLDB*, 11(12):2140–2149, 2018. doi: 10.14778/3229863.3240492.
- [6] M. Cannavicchio, L. Ariemma, D. Barbosa, and P. Merialdo. Leveraging wikipedia table schemas for knowledge graph augmentation. In *WebDB*, pages 5:1–5:6, 2018. doi: 10.1145/3201463.3201468.
- [7] H.-H. Chen, S.-C. Tsai, and J.-H. Tsai. Mining tables from large scale HTML texts. In *COLING*, pages 166–172, 2000. URL <http://aclweb.org/anthology/C00-1025>.
- [8] P. Christen. *Data Matching - Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*. Springer, 2012. doi: 10.1007/978-3-642-31164-2.
- [9] X. Chu, Y. He, K. Chakrabarti, and K. Ganjam. TEGRA: table extraction by global record alignment. In *SIGMOD*, pages 1713–1728, 2015. doi: 10.1145/2723372.2723725.
- [10] X. Chu, J. Morcos, I. F. Ilyas, M. Ouzzani, P. Papotti, N. Tang, and Y. Ye. KATARA: a data cleaning system powered by knowledge bases and crowdsourcing. In *SIGMOD Conference*, pages 1247–1261, 2015. doi: 10.1145/2723372.2749431.
- [11] A. Cimmino and R. Corchuelo. On feeding business systems with linked resources from the Web of Data. In *BIS*, pages 307–320, 2018. doi: 10.1007/978-3-319-93931-5\_22.
- [12] A. Cimmino and R. Corchuelo. A hybrid genetic-bootstrapping approach to link resources in the Web of Data. In *HAIS*, pages 145–157, 2018. doi: 10.1007/978-3-319-92639-1\_13.
- [13] W. W. Cohen, M. Hurst, and L. S. Jensen. A flexible learning system for wrapping tables and lists in HTML documents. In *WWW*, pages 232–241, 2002. doi: 10.1145/511446.511477.
- [14] A. Costa-Silva, A. M. Jorge, and L. Torgo. Design of an end-to-end method to extract information from tables. *IJDAR*, 8(2-3):144–171, 2006. doi: 10.1007/s10032-005-0001-x.

- [15] E. Crestan and P. Pantel. A fine-grained taxonomy of tables on the Web. In *CIKM*, pages 1405–1408, 2010. doi: 10.1145/1871437.1871633.
- [16] E. Crestan and P. Pantel. Web-scale table census and classification. In *WSDM*, pages 545–554, 2011. doi: 10.1145/1935826.1935904.
- [17] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang. Knowledge vault: a web-scale approach to probabilistic knowledge fusion. In *KDD*, pages 601–610, 2014. doi: 10.1145/2623330.2623623.
- [18] J. Eberius, K. Braunschweig, M. Hentsch, M. Thiele, A. Ahmadov, and W. Lehner. Building the Dresden Web Table Corpus: a classification approach. In *BDC*, pages 41–50, 2015. doi: 10.1109/BDC.2015.30.
- [19] V. Efthymiou, O. Hassanzadeh, M. Rodríguez-Muro, and V. Christophides. Matching web tables with knowledge base entities: from entity lookups to entity embeddings. In *ISWC*, pages 260–277, 2017. doi: 10.1007/978-3-319-68288-4\_16.
- [20] H. Elmeleegy, J. Madhavan, and A. Y. Halevy. Harvesting relational tables from lists on the Web. *VLDB*, 20(2):209–226, 2011. doi: 10.1007/s00778-011-0223-0.
- [21] D. W. Embley, M. Hurst, D. P. Lopresti, and G. Nagy. Table-processing paradigms: a research survey. *IJDAR*, 8(2-3):66–86, 2006. doi: 10.1007/s10032-006-0017-x.
- [22] D. W. Embley, M. S. Krishnamoorthy, G. Nagy, and S. C. Seth. Converting heterogeneous statistical tables on the Web to searchable databases. *IJDAR*, 19(2):119–138, 2016. doi: 10.1007/s10032-016-0259-1.
- [23] J. Fan, M. Lu, B. C. Ooi, W.-C. Tan, and M. Zhang. A hybrid machine-crowdsourcing system for matching web tables. In *ICDE*, pages 976–987, 2014. doi: 10.1109/ICDE.2014.6816716.
- [24] F. Fumarola, T. Weninger, R. Barber, D. Malerba, and J. Han. Extracting general lists from web documents: a hybrid approach. In *IEAAIE*, pages 285–294, 2011. doi: 10.1007/978-3-642-21822-4\_29.
- [25] M. Galkin, D. Mouromtsev, and S. Auer. Identifying web tables: supporting a neglected type of content on the Web. In *KESW*, pages 48–62, 2015. doi: 10.1007/978-3-319-24543-0\_4.
- [26] W. Gatterbauer, P. Bohunsky, M. Herzog, B. Krüpl, and B. Pollak. Towards domain-independent information extraction from web tables. In *WWW*, pages 71–80, 2007. doi: 10.1145/1242572.1242583.
- [27] M. Hurst. Layout and language: challenges for table understanding on the Web. In *WDA*, pages 27–30, 2001. URL [http://wda2001.csc.liv.ac.uk/Papers/12\\_hurst\\_wda2001.pdf](http://wda2001.csc.liv.ac.uk/Papers/12_hurst_wda2001.pdf).

- [28] M. Hurst. Classifying TABLE elements in HTML. In *WWW*, 2002. URL <http://wwwconference.org/proceedings/www2002/poster/115/index.html>.
- [29] P. Jiménez, R. Corchuelo, and H. A. Sleiman. ARIEX: automated ranking of information extractors. *Knowl.-Based Syst.*, 93:84–108, 2016. doi: 10.1016/j.knosys.2015.11.004.
- [30] S.-W. Jung and H.-C. Kwon. A scalable hybrid approach for extracting head components from web tables. *IEEE Trans. Knowl. Data Eng.*, 18(2): 174–187, 2006. doi: 10.1109/TKDE.2006.19.
- [31] Z. Khayyat, I. F. Ilyas, A. Jindal, S. Madden, M. Ouzzani, P. Papotti, J. Quiané-Ruiz, N. Tang, and S. Yin. BigDancing: a system for big data cleansing. In *SIGMOD Conference*, pages 1215–1230, 2015. doi: 10.1145/2723372.2747646.
- [32] Y.-S. Kim and K.-H. Lee. Detecting tables in web documents. *Eng. Appl. of AI*, 18(6):745–757, 2005. doi: 10.1016/j.engappai.2005.01.009.
- [33] C. A. Knoblock, P. A. Szekely, E. E. Fink, D. Degler, D. Newbury, R. Sanderson, K. Blanch, S. Snyder, N. Chheda, N. Jain, R. R. Krishna, N. B. Sreekanth, and Y. Yao. Lessons learned in building linked data for the American Art Collaborative. In *ISWC*, pages 263–279, 2017. doi: 10.1007/978-3-319-68204-4\_26.
- [34] L. R. Lautert, M. M. Scheidt, and C. F. Dorneles. Web table taxonomy and formalization. *SIGMOD Record*, 42(3):28–33, 2013. doi: 10.1145/2536669.2536674.
- [35] O. Lehmberg, D. Ritze, R. Meusel, and C. Bizer. A large public corpus of web tables containing time and context meta-data. In *WWW*, pages 75–76, 2016. doi: 10.1145/2872518.2889386.
- [36] K. Lerman, C. Knoblock, and S. Minton. Automatic data extraction from lists and tables in web sources. In *IJCAI*, 2001. URL <http://www.isi.edu/integration/papers/lerman01-atem.pdf>.
- [37] K. Lerman, L. Getoor, S. Minton, and C. A. Knoblock. Using the structure of web sites for automatic segmentation of tables. In *SIGMOD*, pages 119–130, 2004. doi: 10.1145/1007568.1007584.
- [38] T. Liao, T. Liu, S. Zhang, and Z. Liu. Research on web table positioning technology based on table structure and heuristic rules. In *AISC*, pages 351–360, 2018. doi: 10.1007/978-3-319-67071-3\_41.
- [39] X. Ling, A. Y. Halevy, F. Wu, and C. Yu. Synthesizing union tables from the Web. In *IJCAI*, 2013. URL <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6758>.



- [40] M.-L. Lo, K.-L. Wu, and P. S. Yu. TabSum: a flexible and dynamic table summarization approach. In *ICDCS*, pages 628–635, 2000. doi: 10.1109/ICDCS.2000.840979.
- [41] D. P. Lopresti and G. Nagy. Automated table processing: an (opinionated) survey. In *GREC*, pages 109–134, 1999. URL <http://www.cse.lehigh.edu/~lopresti/Publications/1999/grec99.pdf>.
- [42] D. P. Lopresti and G. Nagy. A tabular survey of automated table processing. In *GREC*, pages 93–120, 2000. doi: 10.1007/3-540-40953-X\_9.
- [43] J. Mankoff, H. Fait, and T. Tran. Is your web page accessible? A comparative study of methods for assessing web page accessibility for the blind. In *CHI*, pages 41–50, 2005. doi: 10.1145/1054972.1054979.
- [44] N. Milošević, C. Gregson, R. Hernandez, and G. Nenadic. Disentangling the structure of tables in scientific literature. In *NLDB*, pages 162–174, 2016. doi: 10.1007/978-3-319-41754-7\_14.
- [45] V. Mulwad, T. Finin, Z. Syed, and A. Joshi. Using Linked Data to interpret tables. In *COLD*, 2010. URL [http://ceur-ws.org/Vol-665/MulwadEtAl\\_COLD2010.pdf](http://ceur-ws.org/Vol-665/MulwadEtAl_COLD2010.pdf).
- [46] K. Nishida, K. Sadamitsu, R. Higashinaka, and Y. Matsuo. Understanding the semantic structures of tables with a hybrid deep neural network architecture. In *AAAI*, pages 168–174, 2017. URL <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14396>.
- [47] H. Okada and T. Miura. Detection of layout-purpose table tags based on machine learning. In *UAHCI*, pages 116–123, 2007. doi: 10.1007/978-3-540-73283-9\_14.
- [48] R. K. Padmanabhan, R. C. Jandhyala, M. S. Krishnamoorthy, G. Nagy, S. C. Seth, and W. Silversmith. Interactive conversion of web tables. In *GREC*, pages 25–36, 2009. doi: 10.1007/978-3-642-13728-0\_3.
- [49] G. Penn, J. Hu, H. Luo, and R. T. McDonald. Flexible web document analysis for delivery to narrow-bandwidth devices. In *ICDAR*, pages 1074–1078, 2001. doi: 10.1109/ICDAR.2001.953951.
- [50] C. Peterson. *Learning Responsive Web Design*. O’Reilly, 2014.
- [51] R. Pimplikar and S. Sarawagi. Answering table queries on the Web using column keywords. *PVLDB*, 5(10):908–919, 2012. doi: 10.14778/2336664.2336665.
- [52] F. Qi, X. Wu, and N. Wang. Building top- $k$  consistent results for web table augmentation. In *WISA*, pages 74–79, 2017. doi: 10.1109/WISA.2017.30.

- [53] L.-A. Ratinov, D. Roth, D. Downey, and M. Anderson. Local and global algorithms for disambiguation to Wikipedia. In *ACL*, pages 1375–1384, 2011. URL <http://www.aclweb.org/anthology/P11-1138>.
- [54] X. Ren, Z. Wu, W. He, M. Qu, C. R. Voss, H. Ji, T. F. Abdelzaher, and J. Han. CoType: joint extraction of typed entities and relations with knowledge bases. In *WWW*, pages 1015–1024, 2017. doi: 10.1145/3038912.3052708.
- [55] D. Ritze and C. Bizer. Matching web tables to DBpedia: a feature utility study. In *EDBT*, pages 210–221, 2017. doi: 10.5441/002/edbt.2017.20.
- [56] A. D. Sarma, L. Fang, N. Gupta, A. Y. Halevy, H. Lee, F. Wu, R. Xin, and C. Yu. Finding related tables. In *SIGMOD*, pages 817–828, 2012. doi: 10.1145/2213836.2213962.
- [57] Y. A. Sekhavat, F. D. Paolo, D. Barbosa, and P. Merialdo. Knowledge base augmentation using tabular data. In *LDOW*, 2014. URL [http://ceur-ws.org/Vol-1184/ldow2014\\_paper\\_02.pdf](http://ceur-ws.org/Vol-1184/ldow2014_paper_02.pdf).
- [58] J. L. Sierra, A. Fernández-Valmayor, and B. Fernández-Manjón. From documents to applications using markup languages. *IEEE Software*, 25(2): 68–76, 2008. doi: 10.1109/MS.2008.36.
- [59] J. W. Son and S.-B. Park. Web table discrimination with composition of rich structural and content information. *Appl. Soft Comput.*, 13(1):47–57, 2013. doi: 10.1016/j.asoc.2012.07.025.
- [60] M. Taheriyani, C. A. Knoblock, P. A. Szekely, and J. L. Ambite. Learning the semantics of structured data sources. *J. Web Semant.*, 37-38:152–169, 2016. doi: 10.1016/j.websem.2015.12.003.
- [61] I. Taleb, R. Dssouli, and M. A. Serhani. Big Data pre-processing: a quality framework. In *IEEE Intl. Congress on Big Data*, pages 191–198, 2015. doi: 10.1109/BigDataCongress.2015.35.
- [62] F. Tschirschnitz, T. Papenbrock, and F. Naumann. Detecting inclusion dependencies on very many tables. *ACM Trans. Database Syst.*, 42(3): 18:1–18:29, 2017. doi: 10.1145/3105959.
- [63] P. Venetis, A. Y. Halevy, J. Madhavan, M. Pasca, W. Shen, F. Wu, G. Miao, and C. Wu. Recovering semantics of tables on the Web. *PVLDB*, 4(9): 528–538, 2011. doi: 10.14778/2002938.2002939.
- [64] Y. Wang and J. Hu. Detecting tables in HTML documents. In *DAS*, pages 249–260, 2002. doi: 10.1007/3-540-45869-7\_29.
- [65] K.-L. Wu, S.-K. Chen, and P. S. Yu. Dynamic refinement of table summarization for m-commerce. In *WECWIS*, pages 179–186, 2002. doi: 10.1109/WECWIS.2002.1021257.

- [66] X. Wu, C. Cao, Y. Wang, J. Fu, and S. Wang. Extracting knowledge from web tables based on DOM tree similarity. In *KSEM*, pages 302–313, 2016. doi: 10.1007/978-3-319-47650-6\_24.
- [67] M. Yakout, K. Ganjam, K. Chakrabarti, and S. Chaudhuri. Infogather: entity augmentation and attribute discovery by holistic matching with web tables. In *SIGMOD Conference*, pages 97–108, 2012. doi: 10.1145/2213836.2213848.
- [68] Y. Yang and W.-S. Luk. A framework for web table mining. In *WIDM*, pages 36–42, 2002. doi: 10.1145/584931.584940.
- [69] M. Yoshida, K. Torisawa, and J. Tsujii. A method to integrate tables of the World Wide Web. In *WDA*, pages 31–34, 2001. URL [http://wda2001.csc.liv.ac.uk/Papers/13\\_yoshida\\_wda2001.pdf](http://wda2001.csc.liv.ac.uk/Papers/13_yoshida_wda2001.pdf).
- [70] R. Zanibbi, D. Blostein, and J. R. Cordy. A survey of table recognition. *IJDAR*, 7(1):1–16, 2004. doi: 10.1007/s10032-004-0120-9.
- [71] M. Zhang and K. Chakrabarti. InfoGather+: semantic matching and annotation of numeric and time-varying attributes in web tables. In *SIGMOD Conference*, pages 145–156, 2013. doi: 10.1145/2463676.2465276.
- [72] X. Zhang, Y. Chen, J. Chen, X. Du, and L. Zou. Mapping entity-attribute web tables to web-scale knowledge bases. In *DASFAA*, pages 108–122, 2013. doi: 10.1007/978-3-642-37450-0\_8.