

Mining Data Regions from Web Pages

P S Hiremath
Dept of Computer Science
Gulbarga University
Gulbarga
Karnataka ,India
hiremathps@yahoo.co.in

S S Benchalli
Dept of Computer Science
Karnataka University
Dharwad
Karnataka, India
benchalli_math@yahoo.com

Siddu P Algur
Dept of Info Sc & Engg
SDM College of Engg
& Tech, Dharwad
Karnataka ,India
siddu_p_algur@hotmail.com

Renuka V Udupudi
Dept of Computer
Network Engg ,
V.T.U. , Belgaum
Karnataka, India
renucka@rediffmail.com

Abstract

With the explosive growth of information sources available on the World Wide Web, it has become increasingly difficult to identify the relevant pieces of information, since web pages are often cluttered with irrelevant content like advertisements, navigation-panels, copyright notices etc., surrounding the main content of the web page. A significant amount of valuable information on the web is generated from databases. Pages in data-intensive web sites are usually automatically generated from the back-end DBMS using scripts [6]. Hence, it is useful to mine such data regions in order to extract information from them to provide value-added services. In order to compare the prices and features for products from the various web sites, we need tools to extract attribute descriptions of each product (called *data object*) within a specific region (called *data region*) in pages. In order to accomplish this task, we must identify data of interest by mining relevant data regions from web pages. Currently available automatic techniques [10,14,15] to mine data regions from web pages are still unsatisfactory because of their poor performance. In this paper, we propose a much more effective automatic technique to perform the task. This technique is based on three important observations about data regions on the web. The proposed technique is able to mine the non-contiguous data records and can correctly identify data regions, irrespective of the type of tag in which it is bound. Existing techniques are tag-dependant and work only for records bound within specific tags like <TABLE>, <TR>, <TD>, etc. Our experimental results show that the proposed technique outperforms existing techniques substantially

Categories and Subject Descriptors

Data Mining and Knowledge Discovery.
Web information, web services and semantic web.

Keywords

Web mining, Web data regions, Web data records

1. Introduction

Web information extraction is an important problem for information integration, because multiple web pages may present the same or similar information using completely different formats or syntaxes, which makes integration of information a challenging task. Due to the heterogeneity and lack of structure of web data, automated discovery of targeted information becomes a complex task. A typical web page consists of many blocks or areas, e.g., main content areas, navigation areas, advertisements, etc.

For a particular application, only part of the information is useful, and the rest are noises. Hence it is useful to separate these areas automatically for several practical applications.

Pages in data-intensive web sites are usually automatically generated from the back-end DBMS using scripts [6]. Hence, the structured data on the web are often very important since they represent their host page's essential information, e.g., details about the list of products and services.

In order to extract and make use of information from multiple sites to provide value added services, one needs to semantically integrate information from multiple sources. This problem has been studied by researchers in AI, database and data mining, and web communities [9]. There are several approaches [2,4,5,6,14] for structured data extraction, which is also called wrapper generation.

The first approach [2] is to manually write an extraction program for each web site based on observed format patterns of the site. This manual approach is very labor intensive and time consuming. It thus does not scale to a large number of sites.

The second approach [4] is wrapper induction or wrapper learning, which is currently the main technique. Wrapper learning works as follows: The user first manually labels a set of trained pages. A learning system then generates rules from the training pages. The resulting rules are then applied to extract target items from web pages. These methods either require prior syntactic knowledge or substantial manual efforts. Example wrapper induction systems include WEIN [1]. The third approach [5] is the automatic approach. Since structured data objects on the web are normally database records retrieved from underlying web databases and displayed in web pages with some fixed templates. Automatic methods aim to find patterns/grammars from the web pages and then use them to extract data. Examples of automatic systems are IEPAD [5], ROADRUNNER [6]. These make use of the Patricia (PAT) tree for discovering the record boundaries automatically and a pattern-based extraction rule to extract the web data. This method has a poor performance due to the various limitations of the PAT tree. ROADRUNNER [6] extracts a template by analyzing a pair of web pages of the same class at a time. It uses one page to derive an initial template and then tries to match the second page with the template. Deriving of the initial template has to be again done manually, which is a major limitation of this approach. Another problem with the existing automatic approaches is their assumption that the relevant information of a data record is contained in a contiguous segment of HTML code, which is not always true. MDR [10] basically exploits the regularities in the HTML tag structure directly. It is often very difficult to derive accurate wrappers entirely based on HTML tags. MDR algorithm makes use of the HTML tag tree of the web page to extract data records from the page. However, an incorrect tag tree may be constructed due to the misuse of HTML tags, which in turn makes it impossible to extract data records correctly. MDR has several other limitations which will be discussed in the latter half of this paper. DEPTA [15] uses visual information (locations on the screen at which the tags are rendered) to find data records. Rather than analyzing the HTML code, the visual information is utilized to infer the structural relationship among tags and to construct a tag tree. But this method of constructing a tag tree has the limitation that , the tag tree can be built correctly only as long as the browser is able to render the page correctly. The computation time for constructing the tag tree is also an overhead. However, this method also fails to identify some of the data records. VIPS [11] is based on certain visual cues within the content of a node. Yu et al. [11] suggest that tags such as <HR>, which is used to create a horizontal line, as well as attributes that indicate a change in background color can indicate the beginning

or end of a segment. These segmentation cues are seen as "visual separators" and are classified as horizontal or vertical lines. After the VIPS algorithm has parsed the HTML code, visual separators are detected in the parse tree. The separators receive weights which are adjusted depending on constraints based on separator dimension, adjacent tags such as <HR> (a line), and surrounding font sizes. The separators divide sub-trees further, depending on if they cross a "visual" block of information or not. Finally, the content structure of the page is created, by merging "visual" blocks that are not divided by separators. Knowledge on document structure is used to enhance the quality of query expansion terms in information retrieval. VIPS also makes use of a number of heuristics that do not hold good for all web pages. VIPS also has certain limitations. It does not correctly identify the data regions and it is also dependent on number of heuristic rules which cannot be applied for most of the pages.

We propose a novel and more effective method to mine the data region in a web page automatically. The algorithm is called VSAP (Visual Structure based Analysis of web Pages). It finds the data regions formed by all types of tags using visual cues.

The VSAP technique is based on the observations which are found to be true in the experiments. The experimental results show that the proposed technique outperforms existing techniques substantially.

2. Related Work

Extracting the regularly structured data records from web pages is an important problem. So far, several attempts have been made to deal with the problem.

Related works, mainly in the area of mining data records in a web page automatically, are MDR [10], DEPTA [15].

The first approach, MDR [10] automatically mines all data records formed by table and form related tags i.e, <TABLE>, <FORM>, <TR>, <TD>, etc. assuming that a large majority of web data records are formed by them. The algorithm is based on two observations:

- (a) A group of data records are always presented in a contiguous region of the web page and are formatted using similar HTML tags. Such region is called a *Data Region*.
- (b) The nested structure of the HTML tags in a web page usually forms a tag tree and a set of similar data records are formed by some child sub-trees of the same parent node.

The algorithm works in three steps:

- Step 1. Building the HTML tag tree by following the nested blocks of the HTML tags in the web page.

Step 2. Identifying the data regions by finding the existence of multiple similar generalized nodes of a tag node. A *generalized node* (or a *node combination*) is a collection of child nodes of a tag node, with the following two properties:

- (i) All the nodes have the same parent.
- (ii) The nodes are adjacent.

Then each generalized node is checked to decide if it contains multiple records or only one record. This is done by string comparison of all possible combinations of component nodes using Normalized Edit Distance method [1].

A *data region* is a collection of two or more generalized nodes with the following properties:

- (i) the generalized nodes all have the same parent.
- (ii) the generalized nodes all have the same length.
- (iii) the generalized nodes are all adjacent.
- (iv) the normalized edit distance (string comparison) between adjacent generalized nodes is less than a fixed threshold.

To find relevant data region, MDR makes use of the content mining.

Step 3. Identifying the data records involves finding the data records from each generalized node in a data region.

All the three steps of MDR have certain serious limitations which will be discussed in the latter half of the paper.

The second approach, DEPTA (Data Extraction based on Partial Tree Alignment), also known as MDR-2, is similar to MDR in its data region identification mechanism. DEPTA uses visual information to build the HTML tag tree, to eliminate the problems caused by HTML tag structure irregularities. DEPTA uses tree edit distance for identifying similar data records. A single data record may be composed of multiple sub-trees due to noisy information. MDR may find wrong combinations of sub-trees. DEPTA [15] makes use of the visual gaps between data records to deal with the above problem.

However, as in [10, 15] it also requires content analysis to identify the main data region from the set of data regions obtained. Further, it also relies on TABLE tags for identifying data regions.

This technique proposes a two-step strategy.

- Step 1. Identifying individual data records in a page.
- Step 2. Aligning and extracting data items from the identified data records.

Specifically, this method also uses visual cues to find data records. Since DEPTA is tag dependent, it requires considerable time in building tag tree, traversing whole tag tree and string comparison.

The above automatic methods are inaccurate, tag dependant, incorporate time-consuming tag tree construction, and make many assumptions which do

not always hold good for all web pages. The proposed method does not make any such assumptions and can scale well for almost all web pages. It is also independent of the type of tags and eliminates the time-consuming tag tree construction procedure.

3. The Proposed Technique

We propose a novel and more effective method to mine the data region in a web page automatically. The algorithm is called VSAP (Visual Structure based Analysis of web Pages). The visual information (i.e, the locations on the screen at which tags are rendered) helps the system in three ways:

- (a) It enables the system to identify gaps that separate data records, which helps to segment data records correctly, because the gap within a data record (if any) is typically smaller than that in between data records.
- (b) The visual or display information also contains information about the hierarchical structure of the tags.
- (c) By the visual structure analysis of the web pages, it can be analyzed that the relevant data region seems to occupy the major central portion of the web page.

The system model of the VSAP technique is shown in Fig 1.

It consists of the following components.

- Parsing and Rendering Engine.
- Largest Rectangle Identifier
- Container Identifier
- Data Region Identifier

The output of each component is the input for the next component.

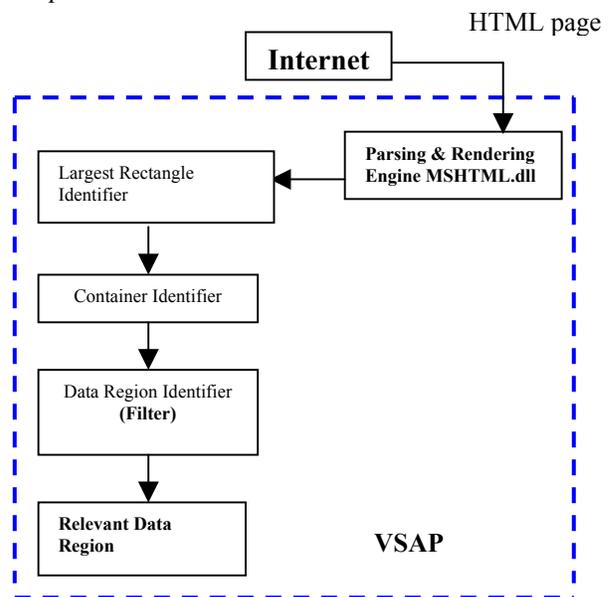


Fig 1. System Model

The VSAP technique is based on three *observations*:

- (a) A group of data records, that contains descriptions of a set of similar objects, is typically presented in a contiguous region of a page.
- (b) The area covered by a rectangle that bounds the data region is more than the area covered by rectangles bounding other regions, eg. advertisements and links.
- (c) The height of a irrelevant data record within a collection of data records is less than the average height of relevant data records within that region.

The experiments show that these observation are true.

Definition 1: A *data region* is defined as the most relevant portion of a web page.

e.g. A region on the product-related web-site that contains a list of products forms the *data region*.

Definition 2: A *data record* is defined as a collection of data that together represents a meaningful independent entity.

e.g. A product listed inside a data region on a product-related website is a data record

Fig 2 illustrates an example, which is a segment of a web page that shows a data region containing list of four books. The full description of each book is a data record.



Fig. 2 An example of a Data Region containing 4 data records

The overall algorithm of the proposed technique is as follows :

Algorithm VSAP(*HTML document*)

- 1 Set maxRect=NULL
- 2 Set dataRegion=NULL
- 3 FindMaxRect(*BODY*);
- 4 FindDataRegion(*maxRect*);
- 5 FilterDataRegion(*dataRegion*);

end

The lines 1 and 2 specify initializations. The line 3 finds the largest rectangle within the container. Line 4 identifies the data region which consists of the relevant data region and some irrelevant regions also. Line 5 identifies the actual relevant data region by filtering the bounding irrelevant regions.

As mentioned earlier, the proposed technique has two main steps. This section presents them in turn.

3.1 Determining the co-ordinates of all bounding rectangles

In the first step of the proposed technique, we determine the co-ordinates of all the bounding rectangles in the web page. The VSAP approach uses the MSHTML parsing and rendering engine of Microsoft Internet Explorer 6.0. This parsing and rendering engine of the web browser gives us these co-ordinates of a bounding rectangle.

The MSHTML parsing and rendering engine is the main HTML component of the Microsoft Internet Explorer web browser. The IE rendering engine is nothing more than a COM component. The rendering engine of the browser produces the boundary coordinates.

We scan the HTML file for tags. For each tag encountered, we determine the co-ordinate of the top-left corner, height and width of the *bounding rectangle* of that tag.

Definition: Every HTML tag specifies a method for rendering the information contained within it. For each tag, there exists an associated rectangular area on the screen. Any information contained within this rectangular area obeys the rendering rules associated with the tag. This rectangle is called the *bounding rectangle* for the particular tag.

A bounding rectangle is constructed by obtaining the co-ordinate of the top-left corner of the tag, the height and the width of that tag. The left and top co-ordinates of the tag are obtained from the offsetLeft and offsetTop properties of the HTMLObjectElement. These values are with respect to it's parent tag. The height and width of that tag are available from the offsetHeight and offsetWidth properties of the HTMLObjectElement Class.

The Fig. 3 shows a sample web page of a product-related website, which contains the list of books and their descriptions which form the data records inside the data region.



Fig. 3 A Sample Web Page of a product related web-site

For each HTML tag on the web page, there exists an associated rectangular area on the screen, which forms the bounding rectangle for that specific tag. Fig 4 shows the bounding rectangles for the `<TD>` tags of the web page shown in Fig 3.

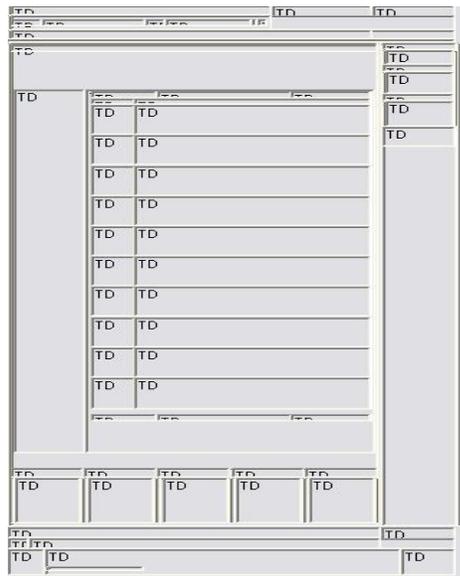


Fig. 4 Bounding rectangles for `<TD>` tag corresponding to the web page in Fig. 3

Figure 5 shows the top-left corner's co-ordinate, height and width obtained for constructing the bounding rectangle of a `<TD>` tag. The bounding rectangle of that tag is shaded. The co-ordinates (80, 45), specify the corresponding height and width of the bounding rectangle of the `<TD>` tag, respectively.

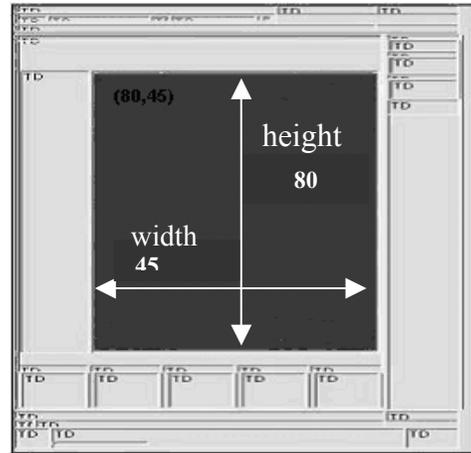


Fig. 5 Bounding rectangle of `<TD>` tag

3.2 Identifying the data regions

The second step of the proposed technique is to identify the data region of the web page. The data region is the most relevant portion of a web page that contains a list of data records.

The three steps involved in identifying the data region are:

- Identify the largest rectangle.
- Identify the container within the largest rectangle.
- Identify the data region containing the data records within this container

3.2.1 Identification of the largest rectangle

Based on the height and width of bounding rectangles obtained in the previous step, we determine the area of the bounding rectangles of each of the children of the BODY tag. We then determine the largest rectangle amongst these bounding rectangles. The reason for doing this is due to the observation that the largest bounding rectangle will always contain the most relevant data in that web page. Thus by determining the largest rectangle, we can thereby obtain a superset of the data region. In Fig 6 the largest rectangle is being shown with a dotted border.

The procedure *FindMaxRect* identifies the largest rectangle amongst all the bounding rectangles of the children of the BODY tag. It is as follows.

Procedure FindMaxRect (BODY)

for each child of BODY tag

Begin

find the co-ordinates of the bounding rectangle for the child

if the area of the bounding rectangle > area of maxRect **then**
maxRect = child

endif

end

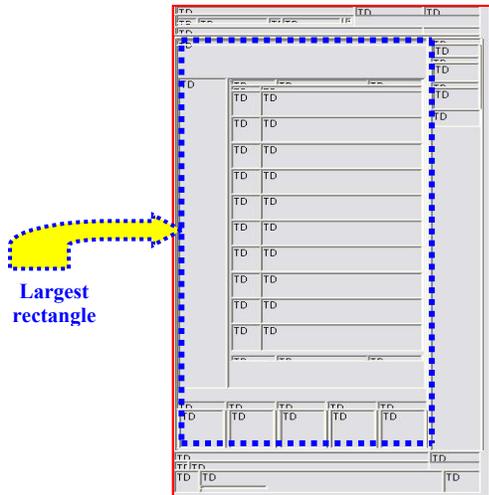


Fig. 6 Largest Rectangle amongst bounding rectangles of Children of BODY tag

3.2.2 Identification of the container within the largest rectangle.

Once we have obtained the largest rectangle, we form a set of all the bounding rectangles whose area is more than half the area of the largest rectangle. The rationale behind this is that the most important data of a web page must occupy a significant portion of the web page. Next, we determine the bounding rectangle having the smallest area in this set. The reason for determining the smallest rectangle within this set is that the smallest rectangle will only contain data records. Thus a *container* is obtained. It contains the data region and some irrelevant data.

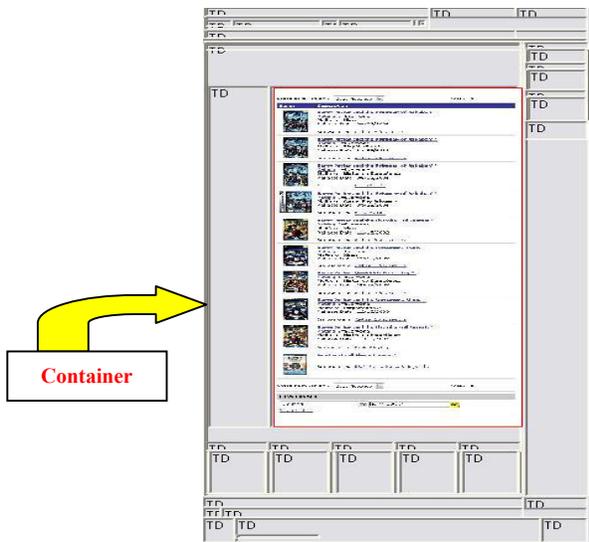


Fig. 7 The container identified from sample web page in Fig. 3

Definition: A *container* is a superset of the data region which may or may not contain irrelevant data.

For example, the irrelevant data contained in the container may include advertisements on the right and bottom of the page and the links on the left side.

The Fig 7 shows the container identified from the web page shown in Fig 3.

The procedure *FindDataRegion* identifies the container in the web page which contains the relevant data region along with some irrelevant data also. It is as follows:

Procedure FindDataRegion (*maxRect*)

ListChildren=depth first listing of the children of the tag associated with *maxRect*
for each tag in ListChildren

Begin

If area of bounding rectangle of tag > half the area of *maxRect* then

If area of bounding rectangle
dataRegion > area of bounding
rectangle of tag then
dataRegion = tag

endif

endif

end

The Fig 8 shows the enlarged view of the container shown in the Fig 7. We note that there is some irrelevant data, both on the top as well as the bottom of the actual data region containing the data records.



Fig. 8 The Enlarged view of the Container shown in Fig. 7

3.2.3 Identification of data region containing data records within the container

To filter the irrelevant data from the container, we use a *filter*. The filter determines the average heights of

children within the container. Those children whose heights are less than the average height are identified as irrelevant data and are filtered off. The Fig 9 shows a filter applied on the container in Fig 8, in order to obtain the data region. We note that the irrelevant data in this case is on the top and bottom of the container, which are being removed by the filter.

The procedure FilterDataRegion filters the irrelevant data from the container, and gives the actual data region as the output. It is as follows:

```

Procedure FilterDataRegion (dataRegion)
    totalHeight = 0
    for each child of dataRegion
        totalHeight += height of the bounding
        rectangle of child
    end
    avgHeight = totalHeight / no of children of
    dataRegion
    for each child of dataRegion
        Begin
        If height of child's bounding
        rectangle < avgHeight then
            Remove child from dataRegion
        endif
    end

```



Fig. 9 Data Region obtained after filtering the Container in Fig. 8

The VSAP technique, as described above, is able to mine the relevant data region containing data records from the given web page efficiently. The VSAP algorithm has been implemented using Visual Basic 6.0 on Windows XP environment..

4. Empirical Evaluation

In this section, we evaluate the proposed technique. We also compare it with two state-of-the-art existing systems, DEPTA [15] (which is an improvement of [10]) and MDR [10]. We do not compare it with the

method in [5] and the method in [6] here as it is shown in [10] that MDR is already more effective than them.

The evaluation consists of three aspects as discussed in the following:

1. Data Region Extraction:

We compare the first step of [10] with our system for identifying the data regions.

Both MDR [10] and DEPTA [15], are dependent on certain tags like <TABLE>, <TBODY>,etc. for identifying the data region. But, a data region need not be always contained only within specific tags like <TABLE>, <TBODY>, etc. A data region may also be contained within tags other than table-related tags like <P>, , <FORMS> etc.

In the proposed VSAP system, the data region identification is independent of specific tags and forms. Unlike [10], where an incorrect tag tree may be constructed due to the misuse of HTML tags, there is no such possibility of erroneous tag tree construction in case of VSAP, because the hierarchy of tags is constructed based on the visual cues on the web page.

In case of [10] and [15] , the entire tag tree needs to be scanned in order to mine data regions, but VSAP does'nt scan the entire tag tree, but it only scans the largest child of the <BODY> tag. Hence, this method proves very efficient in improving the time complexity compared to other contemporary algorithms.

2. Data Record Extraction:

We compare the record extraction step of MDR [10] with VSAP:

MDR identifies the data records based on keyword search (e.g , “ \$ “). But VSAP is purely dependent on the visual structure of the web page only. It does not make use of any text or content mining. This proves to be very advantageous as it overcomes the additional overhead of performing keyword search on the web page.

MDR, not only identifies the relevant data region containing the search result records but also extracts records from all the other sections of the page, e.g, some advertisement records also, which are irrelevant.

MDR also sets a similarity threshold value of 0.3 as the default value for the system to identify similar data records, which has been set based upon a number of training pages. Hence, it does'nt scale well with all web pages.

In MDR, comparison of generalized nodes is based on string comparison using normalized edit distance method. In DEPTA, comparison of sub-trees is made by simple tree matching algorithm. However, both these methods are slow and inefficient as compared to VSAP where the comparison is purely

numeric, since we are comparing the co-ordinates of the two bounding rectangles. It scales well with all the web pages.

A single data record may be composed of multiple sub-trees. Due to noisy information, MDR may find wrong combinations of sub-trees. In VSAP system, visual gaps between data records help to deal with this problem.

3. Overall Time Complexity

Complexity of VSAP is much lesser than the existing algorithms. The existing algorithm MDR [10] has time complexity of the order $O(NK)$ without considering string comparison, where N is the total number of nodes in the tag tree and K is the maximum number of tag nodes that a generalized node can have (which is normally a small number <10). DEPTA [15] has the time complexity of the order $O(k^2)$ without considering tree matching, where k is the number of trees. Our algorithm VSAP has a complexity of the order of $O(n)$, where n is the number of tag-comparisons made. Although, the complexities of these algorithms are expressed in different parameters, the comparison of dominant operations in each of them leads to the conclusion that VSAP is a time saving algorithm.

5. Experimental Results:

In running MDR, we used their default settings. MDR system was downloaded at <http://www.cs.uic.edu/~liub/MDR/MDR-download.htm>. We use the *recall* and *precision* measures (which are widely used to evaluate information retrieval systems) to evaluate the performance of VSAP system for extracting data records. Recall and precision are defined below:

$$Recall = E_c / N_t \quad \text{and} \quad Precision = E_c / E_t$$

where, E_c is the total number of correctly extracted records, N_t is the total number of records on the page, and E_t is the total number of records extracted. *Recall* defines the correctness of the data records identified and *precision* is the percentage of the relevant data records identified from the web page.

The experimental results are given in Table 1. Correct results (Cor.) are relevant data records present on the page that have been correctly identified. Wrong results (Wr.) are irrelevant data records that have been incorrectly identified. The results obtained after running both MDR and VSAP are shown in table 1. The successes of both algorithms are compared in terms of precision and recall.

Column 1 : It lists the URL of each site. In some sites more than one page are tried (which have different data record formats). All our experimental web pages are

selected randomly. Due to long URLs of most pages, we could not list them here.

URL	MDR		VSAP	
	Cor.	Wr.	Cor.	Wr.
http://www.tigerdirect.com/....	8	36/0	8	0/0
http://www.amazon.com/....	0	17/25	25	1/0
http://www.cooking.com/....	17	13/3	20	0/0
http://www.ebay.com/.....	25	30/0	25	0/0
http://www.powells.com/.....	9	47/1	10	5/0
http://www.barnesandnoble.com/....	10	30/0	10	0/0
http://www.pricegrabber.com/.....	0	0/25	25	0/0
http://www.shoebuy.com/....	12	12/84	96	0/0
http://www.smartbuy.com/.....	0	15/10	10	0/0
http://www.reviews.cnet.com/.....	24	38/1	25	3/0
http://www.nothingbutsoftware.com/	10	12/0	10	0/0
http://www.refurbdepot.com/....	0	6/15	15	0/0
http://www.drugstore.com/....	15	14/0	15	0/0
http://www.bookpool.com/....	10	7/0	10	0/0
http://www.target.com/.....	0	0/12	12	1/0
Total	140	277 / 176	316	10 / 0
Recall	33.5%		96.93%	
Precision	44.3%		100%	

Table 1 Comparison of MDR and VSAP

Columns 2 and 4 : They give the numbers of correct (Cor.) records extracted by MDR and VSAP from the pages of each site respectively. These data records are the relevant ones (e.g., product lists). They do not include navigation areas, advertisements etc, which may also have regular patterns. MDR cannot handle nested data records (records within records), but VSAP is able to handle such data records also.

Columns 3 and 5: They give the numbers of data records extracted wrongly (Wr.) by MDR and VSAP, from the pages of each site respectively. x/y means that x is the number of extracted results that are wrong, and y is the number of results that are not extracted.

The results shown are obtained without using any form of content search. All the URLs given are obtained from the search functionality in the browser, where the search query for all URLs is “watch” (except for test case 12 where “camera” was specified, test case 14 where “computer” was specified and 8 where “men’s” was specified)

We note that test cases 7 and 15 hanged in case of MDR and hence have been considered as missing the correct data records. It has been experimentally found out that MDR fails to identify a single data record where as VSAP successfully identifies it for the URL, <http://www.nothingbutsoftware.com/SiteSearch.asp?db=12719&query=watches>

The last three rows of Table 1 give the total of each column, the recall and precision of each system. For MDR and VSAP, the recall and precision are computed based on the total number of data records found in all pages and the actual number of data records in these pages. We see that the data record extraction is highly effective. We also observe that VSAP performs significantly better than MDR.

6. Conclusion

In this paper, we have proposed a new approach to extract structured data from web pages. Although the problem has been studied by several researchers, existing techniques are either inaccurate or make many strong assumptions. A novel and effective method VSAP is proposed to mine the data region in a web page automatically. It is a pure visual structure oriented method that can correctly identify the data region. VSAP is independent of errors occurring due to misuse of HTML tags. Most of the current algorithms fail to correctly determine the data region, when the data region consists of only one data record. Also, most of the sites fail in the case where a series of data records is separated by an advertisement, followed again by a single data record. VSAP works correctly for both the above cases. It is able to correctly identify data records, irrespective of the type of tag in which it is bound. Other methods would work for only tags of certain types. The number of comparisons done in VSAP is significantly lesser than other approaches. Further, the comparisons are made on numbers, unlike other methods where strings or trees are compared. Thus VSAP overcomes the drawbacks of existing methods and performs significantly better than these methods.

Scope for future work:

Extraction of the data fields from the data records contained in these mined data regions will be considered in the future work taking also into account the complexities such as the web pages featuring dynamic html, etc. The extracted data can be put in some suitable format and eventually stored back into a relational database. Thus, data extracted from each web-page can then be integrated into a single collection. This collection of data can be further used for various Knowledge Discovery Applications, e.g., making a comparative study of products from various companies, smart shopping, etc.

7. Acknowledgement

The authors are grateful to the referees for their suggestions and useful comments.

8. References

- [1] Baeza Yates, R. Algorithms for string matching: A survey. *ACM SIGIR Forum*, 23(3-4): 34–58, 1989.
- [2] J. Hammer, H. Garcia Molina, J. Cho, and A. Crespo. Extracting semi-structured information from the web. *In Proc. of the Workshop on the Management of Semi-structured Data*, 1997.
- [3] D. Embley, Y. Jiang, and Y. K. Ng. Record-boundary discovery in Web documents. *ACM SIGMOD Conference*, 1999
- [4] Kushmerick, N. Wrapper Induction: Efficiency and Expressiveness. *Artificial Intelligence*, 118:15-68, 2000. [Clustering-based Approach to Integrating Source Query]
- [5] Chang, C-H., Lui, S-L. IEPAD: Information Extraction Based on Pattern Discovery. *WWW-01*, 2001.]
- [6] Crescenzi, V., Mecca, G. and Merialdo, P. ROADRUNNER: Towards Automatic Data Extraction from Large Web Sites. *VLDB-01*, 2001.]
- [7] Eying, H. Zhang. HTML Page Analysis based on Visual Cues. *6th International Conference on Document Analysis and Recognition*, 2001.
- [8] D. Buttler, L. Liu, C. Pu. A Fully Automated Object Extraction System for the World Wide Web. *International Conference on Distributed Computing Systems (ICDCS 2001)*, 2001
- [9] Bing Liu, Kevin chen-chuan chang, Editorial: Special issue on web content mining, *WWW 02*, 2002.
- [10] Liu, B., Grossman, R. and Zhai, Y. Mining Data Records in Web Pages. *KDD-03*, 2003.
- [11] Cai, D., Yu, S., Wen, J.-R. and Ma, W.-Y. (2003). Extracting Content Structure for Web Pages based on Visual Representation, *Asia Pacific Web Conference (APWeb 2003)*, pp. 406-417.
- [12] A. Arasu, H. Garcia-Molina, Extracting structured data from web pages, *ACM SIGMOD 2003*, 2003
- [13] J. Wang, F. H. Lochovsky. Data Extraction and Label Assignment for Web Databases. *WWW conference*, 2003.
- [14] H. Zhao, W. Meng, Z. Wu, Raghavan, Clement Yu. Fully Automatic Wrapper Generation For Search Engines, *International WWW conference 2005, May 10-14, 2005, Japan. ACM 1-59593-046-9/05/005*
- [15] Zhai, Y., Liu, B. Web Data Extraction Based on Partial Tree Alignment, *WWW-05, 2005, May 10-14, 2005, Chiba, Japan. ACM 1-59593-046-9/05/00*