

EXPERIMENTING WITH RULE LEARNING FOR INFORMATION EXTRACTION FROM HTML

Presented at 6th Int. Symposium SYNASC04, Timișoara,
Romania

Costin Bădică *, Amelia Bădică **

Abstract. The Web is a continuously growing information repository with a rich semantic structure that spans many application areas. The Web, however, has been designed primarily for human consumption rather than automated processing. This is a major obstacle for automating tasks like information searching, filtering and extraction. In this context, the aim of the paper is to present a technique for learning rules to extract product information from HTML information sources that represent product information sheets. The technique exploits the fact that the Web pages that represent product information of a certain producer are generated on the fly from the producer database and therefore they exhibit uniform structures. Consequently, while the extraction task is executed manually for a few information items by a human user, a general-purpose inductive learner can learn extraction rules that will be further applied to the current and other product information sheets to automatically extract other items. The input to the learning algorithm is a relational description of the HTML document tree that defines the HTML tree nodes types and the relationships between them. The approach is demonstrated with appropriate examples, experimental results, and software tools.

Key words: Web information extraction, first-order inductive learning

AMS Subject Classification: 68P20, 68T35, 68U35

*University of Craiova, Software Engineering Department, Bvd.Decebal 107, Craiova, RO-200440, Romania, badica_costin@software.ucv.ro

**University of Craiova, Business Information Systems Department, A.I.Cuza 13, Craiova, RO-200585, Romania, ameliabd@yahoo.com

1. Introduction. The Web is now a huge information repository that is characterized by i) *high diversity*, i.e. the Web information covers almost any application area, ii) *disparity*, i.e. the Web information comes in many formats ranging from plain and structured text to multimedia documents and iii) *rapid growth*, i.e. old information is continuously being updated in form and content and new information is constantly being produced. The HTML markup language is the *lingua franca* for publishing information on the Web. In our opinion these facts explain the significant growth of the interest in the field of automatizing the tasks of information extraction (IE) from HTML information sources.

IE was developed as part of DARPA's MUC program and originally was concerned with locating specific pieces of information in text documents written in natural language (Lenhart, 1991). The field then expanded rapidly to cover extraction tasks from networked documents including HTML and XML.

Paper (Muslea, 1999) identified two major types of IE: i) IE from unstructured texts that treat the document as a string of characters and ii) IE from (semi-)structured texts that treat the document as a tree.

Traditional IE from strings uses techniques from computational linguistics based on finite-state methodologies. Paper (Chidlovskii, 2003) cites two approaches: i) the *local approach* concerned with enhancing document parsers with a set of extraction patterns and ii) the *global approach* concerned with applying grammatical inference methods to identify the language of the target document.

IE from tree documents exploits the structural information conveyed by the document tree. An important advantage of this approach is that the extracted item can depend on its structural context in a document, while this information is lost in the event the tree document is treated linearly as a string (Kosala, 2002). Two recent IE approaches in this class are: i) *tree automata induction* (Kosala, 2002) and ii) *learning subtree delimiters* (Chidlovskii, 2003).

This paper presents a methodology and experimental results on the application of first order rule learning to extract information from HTML documents structured as trees. The experiments were carried out in the area of collecting product information from the Web, an important task in product data integration (Fensel, 2001, Bădică, 2003). However, the approach is general enough to be useful in other application areas that require IE from tree structured documents including XML and XHTML, as for example in wrapper development for querying bibliographic databases

in digital libraries.

The rest of the paper is structured as follows. Section 2 contains an overview of the stages in the IE process. In section 3 we present the relational representation of HTML document trees. In section 4 we show how the task of learning IE rules from a tree document can be mapped to FOIL (Quinlan, 1995). Section 5 contains a brief description of the software tools we have used in our experiments. Section 6 presents the experiments performed and the results obtained in terms of precision and recall measures. Section 7 cites related work highlighting similarities with and differences from our work. Section 8 contains some conclusions and points to future research.

2. Overview of the Extraction Process. The IE process was structured as a linear sequence of stages:

i) *Crawl or browse the Web and download HTML pages of interest.* This step is quite straightforward. Either a human user is browsing the Web downloading interesting pages or the task of navigating the Web and downloading pages is automatized by means of a crawler component ((Chakrabarti, 2003)). The result of this step is a collection of Web pages that were fetched and stored locally on the user machine.

ii) *Preprocess and convert the document from HTML to XHTML.* In this step the input HTML document is cleaned and converted to a well-formed XML document written in XHTML. The resulted XHTML document is structured as a tree.

iii) *Manually extract a few examples.* In this step the user loads a few XHTML pages from the local repository and performs a few extraction tasks. The result is a set of annotated XHTML documents with the extracted fields.

iv) *Parse the annotated XHTML documents and generate a relational representation as input for the learner.* In this step the annotated XHTML documents are parsed and converted into input for the learning program. The result will contain the target relation, the sets of positive and negative examples and the background relations.

v) *Apply the learning algorithm and obtain the IE rules.* In this step the learning program is executed on the input generated in the previous step. The result is a set of first order IE rules.

vi) *Apply the rules to extract information.* In this step the user applies the learnt IE rules to new XHTML pages in order to automatically extract new information. In principle there are two possibilities to perform this

task: i) the target document is converted to a relational representation in Prolog and the first order IE rules are applied to the Prolog representation or ii) the first order IE rules are mapped to XSLT and the resulted XSLT is applied to the target XHTML document.

3. Relational Representation of Document Trees. An XHTML document is composed of a structural component and a content component.

The structural component consists of the set of document nodes or elements. The document elements are nested into a tree like structure. Each document element has assigned a specific tag from a given finite set of tags Σ . There is a special tag *text* $\in \Sigma$ that designates a text element.

The content component of a document consists of the actual text content in the text elements and the attribute-value pairs attached to the other document elements.

The structure of XHTML documents can be represented as unranked ordered trees, also known as Σ -trees. According to (Neven, 2000), the set \mathcal{T}_Σ of Σ -trees is defined inductively as follows:

- i) if $\sigma \in \Sigma$ then $\sigma \in \mathcal{T}_\Sigma$;
- ii) if $\sigma \in \Sigma$ and $t_1, \dots, t_n \in \mathcal{T}_\Sigma$, $n > 0$, then $\sigma(t_1, \dots, t_n) \in \mathcal{T}_\Sigma$.

Note that there is no a priori bound on the number of children of a node in a Σ -tree, i.e. the tree is *unranked*, and note also that the set of children of a given node in a Σ -tree is ordered, i.e. the tree is *ordered*.

Example 1. Consider the Hewlett Packard's site of electronic products³ and the task of IE from a product information sheet for printers. The printer information is displayed in a two-column table as a set of feature-value pairs. Our task is to extract the names and/or the values of the printer features. This information is usually stored in the leaf elements of the page. Figure 1 displays in the left panel the XHTML tree of a fragment of this document and in the right panel the graphical view of this fragment as a two-column table. Figure 2 displays this XHTML document fragment as a Σ -tree. ■

Recall that our task in this section was to devise a relational representation for an XHTML document tree. We assign a unique identifier (an integer value) to each node of the tree. Let \mathcal{N} be the set of all node identifiers.

³<http://www.hp.com>

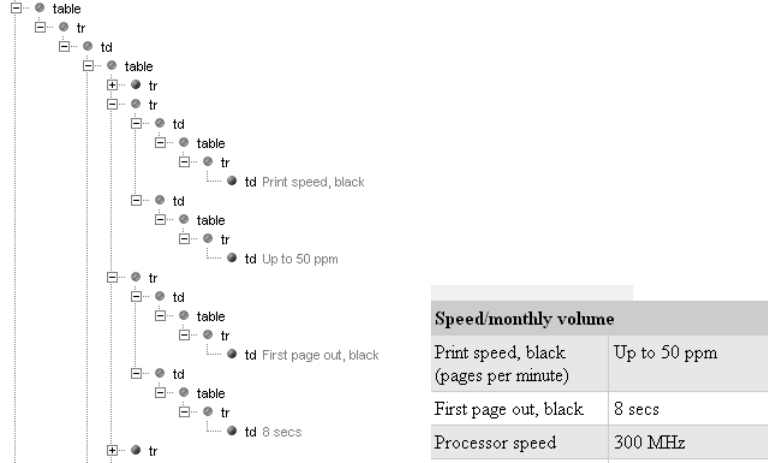


Fig. 1. An XHTML document fragment and its graphic view

The structural component of an XHTML document tree can be represented using the following three relations:

- i) $child \subseteq \mathcal{N} \times \mathcal{N}$, ($child(P, C) = true$) \Leftrightarrow (P is the parent of C).
- ii) $next \subseteq \mathcal{N} \times \mathcal{N}$, ($next(L, N) = true$) \Leftrightarrow (L is the left sibling of N).
- iii) $tag \subseteq \mathcal{N} \times \Sigma$, ($tag(N, T) = true$) \Leftrightarrow (T is the tag of node N).

We introduce two more sets that are needed to represent the content component of an XHTML document tree: the set \mathcal{S} of content elements, which are basically the strings attached to text elements and the values assigned to element attributes, and the set \mathcal{A} of attributes. With these notations, the content component of an XHTML document tree can be represented using the following two relations:

- i) $content \subseteq \mathcal{N} \times \mathcal{S}$, ($content(N, S) = true$) \Leftrightarrow (S is the string contained by the text node N).
- ii) $attribute_value \subseteq \mathcal{N} \times \mathcal{A} \times \mathcal{S}$, ($attribute_value(N, A, S) = true$) \Leftrightarrow (S is the string value of the attribute A of the node N).

Example 2. The right panel of figure 2 displays a part of the relational representation of the XHTML document tree shown on the left panel. ■

The relational representation of document trees introduced in this section makes convenient the definition of first-order IE rules. An IE rule defines a relation $extract \subseteq \mathcal{N}$ such that ($extract(N) = true$) \Leftrightarrow (N is an extracted node).

The definition of an IE rule consists of a set of clauses. The head of a clause is $extract(N)$ and the body of a clause is a conjunction of positive

<i>html</i> (...	<i>tag</i> (0, <i>html</i>)		
<i>table</i> (<i>tag</i> (100, <i>table</i>)		
<tr>(...),</tr>	<i>tag</i> (101, <i>tr</i>)	<i>child</i> (100, 101)	
<tr(< td=""> <td><i>tag</i>(102,<i>tr</i>)</td> <td><i>child</i>(100, 101)</td> </tr(<>	<i>tag</i> (102, <i>tr</i>)	<i>child</i> (100, 101)	
<td><i>table</i>(<i>tr</i>(<i>td</i>(<i>text</i>))),</td> <td><i>tag</i>(103,<i>tr</i>)</td> <td><i>child</i>(100, 102)</td>	<i>table</i> (<i>tr</i> (<i>td</i> (<i>text</i>))),	<i>tag</i> (103, <i>tr</i>)	<i>child</i> (100, 102)
<td><i>td</i>(<i>table</i>(<i>tr</i>(<i>td</i>(<i>text</i>))),</td> <td><i>tag</i>(107,<i>td</i>)</td> <td><i>child</i>(100, 103)</td>	<i>td</i> (<i>table</i> (<i>tr</i> (<i>td</i> (<i>text</i>))),	<i>tag</i> (107, <i>td</i>)	<i>child</i> (100, 103)
),	<i>tag</i> (108, <i>table</i>)	<i>child</i> (101, 107)	
<tr(< td=""> <td><i>tag</i>(109,<i>tr</i>)</td> <td><i>child</i>(107, 108)</td> </tr(<>	<i>tag</i> (109, <i>tr</i>)	<i>child</i> (107, 108)	
<td><i>table</i>(<i>tr</i>(<i>td</i>(<i>text</i>))),</td> <td><i>tag</i>(110,<i>td</i>)</td> <td><i>child</i>(108, 109)</td>	<i>table</i> (<i>tr</i> (<i>td</i> (<i>text</i>))),	<i>tag</i> (110, <i>td</i>)	<i>child</i> (108, 109)
<td><i>td</i>(<i>table</i>(<i>tr</i>(<i>td</i>(<i>text</i>))),</td> <td><i>tag</i>(111,<i>text</i>)</td> <td><i>child</i>(109, 110)</td>	<i>td</i> (<i>table</i> (<i>tr</i> (<i>td</i> (<i>text</i>))),	<i>tag</i> (111, <i>text</i>)	<i>child</i> (109, 110)
),		<i>child</i> (110, 111)	
<tr>(...),</tr>	<i>content</i> (111,		
...	'Print speed,	<i>next</i> (101, 102)	
)	black')	<i>next</i> (102, 103)	
...		<i>next</i> (103, 104)	
)	<i>attr_value</i> (108,		
	<i>border</i> , '0')		

Fig. 2. The structure of the XHTML fragment from figure 1 represented as a Σ -tree and a part of its relational representation.

or negative literals made of the relations *child*, *next*, and *tag* (*attr_value* could be used as well to additionally check node attributes). The relation *content* can be used to extract the actual content of the answer nodes.

Example 3. Assuming that we want to extract all the nodes of the XHTML document from figure 1 that are of type *text* that have a grand-grand-parent of type *table* that has a parent that has a right sibling, we can use the following IE rule:

$$\begin{aligned}
 \text{extract}(A) \leftarrow & \\
 & \text{tag}(A, \text{text}) \wedge \text{child}(B, A) \wedge \text{child}(C, B) \wedge \text{child}(D, C) \wedge \\
 & \text{tag}(D, \text{table}) \wedge \text{child}(E, D) \wedge \text{next}(E, F)
 \end{aligned}$$

■

Obviously, the manual writing of IE rules of this type is a slow and difficult process requiring a careful analysis of the structure of the target XHTML document. Fortunately, it is possible to learn first order IE rules with a general purpose relational learning program and this issue is the focus of the next section of the paper.

4. Using FOIL Program to Learn Extraction Rules. FOIL – First Order Inductive Learner is a general purpose first order (also known as rela-

tional) inductive learning program developed by John Ross Quinlan and his team at the beginning of the '90s ((Quinlan, 1995)). Even if it is not in the scope of the paper to give a detailed description of how FOIL works, we decided to give a brief overview of first order inductive learning in order to make the paper self contained. For more details see (Quinlan, 1995) and (Mitchell, 1997) (chapter 10).

In first order inductive learning the training data comprises the following two components:

i) The *target relation*. This is the relation that is learnt. It is defined extensionally as a set of tuples. Usually the tuples are partitioned into a set of positive tuples and a set of negative tuples.

ii) The *background relations*. They represent the background domain knowledge and are usually defined extensionally as sets of tuples.

The goal of first order inductive learning is to construct a logic program that represents an appropriate intensional definition of the target relation in terms of the background relations and, optionally itself, if recursive definitions are allowed. The learnt definition must cover all the positive tuples (or a large fraction of them) and no negative tuples (or a small fraction of them) from its extensional definition.

FOIL requires as input the information about the target and the background relations, as mentioned above, and the definition of the relations arguments types, as sets of constants. FOIL's output is a set of clauses that represents a disjunctive definition of the target relation. The set of clauses is constructed iteratively, one clause at a time, removing all the positive tuples until an empty set remains. Every clause is constructed using an heuristic hill climbing search strategy that is guided by the information gain, determinate literals and the clause complexity (see (Quinlan, 1995) for details).

An important step in our experiments with IE was to prepare the input for FOIL. In what follows we assume that during the training stage, the user selects a single XHTML document – the training document, and performs a few extraction tasks on it. The problem is to map the training data into input for FOIL.

First, we must define the relations arguments types. We used a single type that represents the set \mathcal{N} of all the nodes found in the training document.

Second, we must define the target relation *extract*. We used only the nodes extracted during the training stage from the training document, as the set of positive examples. FOIL uses the closed world assumption

by considering all the other nodes in the training document as negative examples.

Third, we must define the background relations *next* and *tag*. For convenience, we have replaced the relation *tag* with a set of relations r_σ , for all $\sigma \in \Sigma$, defined as $r_\sigma = \{N | \text{tag}(N, \sigma) = \text{true}\}$.

Example 4. We consider a FOIL input example taken from our experiments. The input is split into the following sections: argument types (N), target relation (*extract*, only positive examples), and background relations (*child*, *next*, *table*, ..., *span*).

#N: 0,1,2,3,4, ..., 1096,1097.

```

extract(N)
1004
1006
1008
...
1056
1058
.
*child(N,N)
0,1
0,2
1,3
1,4
...
1094,1097
.
*next(N,N)
1,2
3,4
...
1091,1092
.
*table(N)
26
27
...
991

```

```

.
...
*span(N)
250
251
...
980
.

```

■

Additionally, FOIL can be parameterized from the command line. All the parameters have default values, but we found useful to control explicitly the following options:

- i) The use or not of negative literals in the body of a clause.
- ii) The maximum variable depth in literals. Variables occurring in the clause head have depth 0 and a new variable in the literal has depth one greater than the maximum depth of its existing variables.
- iii) Finally, to set the minimum clause accuracy to an appropriate value. The clause accuracy represents the percentage of the positive tuples from the set of all tuples covered by the relation. FOIL will not accept any clause with an accuracy lower than this value.

Example 5 This example shows a FOIL command-line that we have used in our experiments.

```
foil -d8 -a40 -v3 <ex10.d >ex10.o
```

The parameter `d8` is used for setting the variable depth to 8. The parameter `a40` is used for setting the minimum clause accuracy to 40 %. The parameter `v3` is used to set the level of verbosity of FOIL's output to 3 (minimum is 0 and maximum is 4), meaning that the output will include quite detailed information about the search process performed by FOIL. The parameters `ex10.d` and `ex10.o` are the files representing FOIL's input and output. ■

5. Software Tools. We have developed a set of prototype software tools to support our experiments with rule learning for IE from HTML. The tool set implements the scenario outlined in section 2 with the restriction that the training stage is based on a single training document. The tool set incorporates the following on-the-shelf software components:

- i) The machine learning program FOIL. This component is called from the main program through a command line.

ii) The Xerces⁴ library for XML document processing. This component is used to read an XHTML document into a DOM tree. The DOM tree is then mapped onto the format required as input for FOIL using a breadth-first traversal.

iii) The Tidy⁵ library for HTML document cleaning and pre-processing. This component is used in the pre-processing stage of HTML documents for converting them to XHTML. We have also developed a custom XSLT stylesheet to remove redundant HTML elements like *script* and *style*.

The development of a Java-based GUI shell for our set of tools is underway. It will support the user in performing the following tasks: convert HTML files to XHTML, load an XHTML document for training, manually select with the mouse some examples for extraction, generate the input for FOIL from the training document and the manually extracted information, run FOIL, visualize the learnt IE rules, convert the IE rules to XSLT, visualize the XSLT IE rules, run the XSLT IE rules on a target XHTML document and visualize the output.

Figure 3 presents our tool in action, during the stage of collecting examples from a training document.

6. Experimental Results. We ran a series of experiments of rule learning for IE from the Hewlett Packard's Web site. The task was to extract the printers feature values from their information sheets.

We selected an experimental set of 30 documents representing information sheets for HP printers. The printer information is represented in two column HTML tables (see figure 1). The feature values are stored in the second column. Each of the selected documents contains between 28 and 38 features.

We used training examples from a single document representing the information sheet of the HP Business Inkjet 2600 (C8109A) printer. This document contains 28 positive examples and a total of 1098 nodes. We performed learning experiments with 1, 2, 3, 4, 12, 24 and 28 positive examples selected from this set. The results of the learning stage were applied to all of the 30 documents, measuring the precision and recall indicators. These values were computed by averaging the precision and recall for each individual document.

⁴<http://xml.apache.org/xerces2-j/index.html>

⁵<http://www.w3c.org/People/Ragget/tidy/>

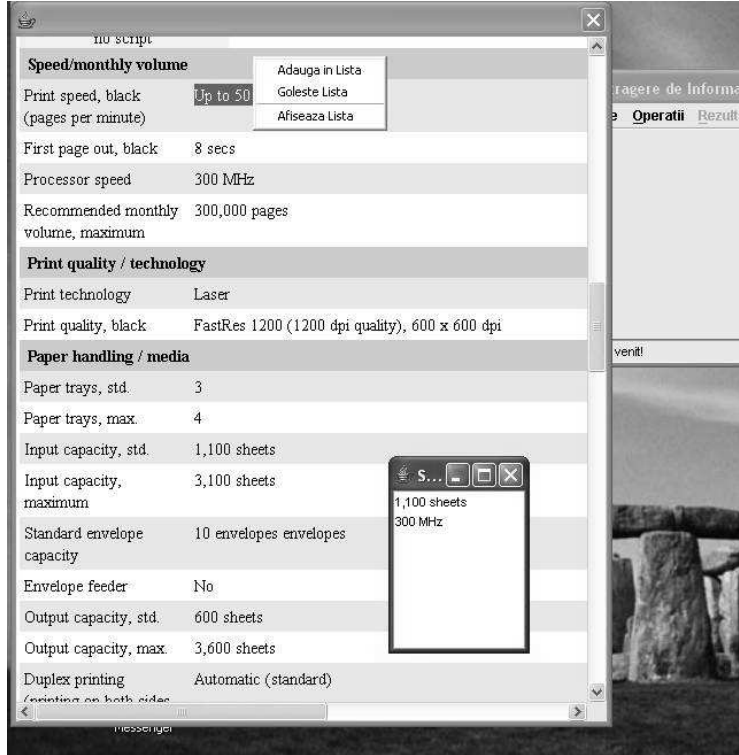


Fig. 3. Software tool for information extraction from HTML.

The experiments were divided in two classes:

- i) In the first class we have ignored the *next* relation as a background relation of FOIL's input.
- ii) In the second class we have added the *next* relation to the set of background relations of FOIL's input.

For each class of experiments we have varied explicitly the following parameters: the use or not of negated literals in rule bodies (FOIL parameter), the minimum clause accuracy (FOIL parameter), and the number of positive examples. We also set the maximum variable depth to 8 and used the closed world assumption in generating the set of negative examples in all the experiments.

The results of the experiments in the first and second class are summarized in tables 1 and 2 respectively.

After a careful analysis of the experimental results of our experiments

Table 1. Experiments without the use of the background relation *next*.

Exper.no.	Neg.lit.	No.pos.ex.	Cl,acc.	Precision	Recall
1	No	12	10 %	0.334	0.980
2	Yes	12	0 %	0.355	1
3	No	24	20 %	0.334	0.980
4	Yes	24	0 %	0.409	0.970
5	No	28	20 %, 30 %	0.334	0.980
6	Yes	28	0 %, 20 %	0.409	0.970

Table 2. Experiments with the use of the background relation *next*

Exper.no.	Neg. lit.	No.pos.ex.	Cl.acc.	Precision	Recall
1	No	1	0 %	0.332	1
2	Yes	1	0 %	0.411	1
3	No	2	0 %	0.730	0.800
4	Yes	2	0 %	0.440	1
5	No	3	0 %, 5 %	0.698	1
6	Yes	3	0 %	0.440	1
7	No	4	0 %	0.332	1
8	No	4	10 %	0.730	0.800
9	Yes	4	0 %	0.448	1
10	No	12	20 %, 40 %	0.855	0.800
11	Yes	12	20 %	1	0.800
12	No	24	20 %	0.698	1
13	No	24	60 %	0.855	0.800
14	Yes	24	80 % (default)	1	0.800
15	No	28	80 % (default)	0.855	0.800
16	Yes	28	80 % (default)	1	1

there are some points to note:

i) In all the experiments with the precision lower than 0.500 the learnt IE rule also extracted the feature name of every extracted feature value. This explains the low values of the precision in these experiments.

ii) The low precision values in the experiments of the first class clearly shows the importance of exploiting the information conveyed by the *next* relation for IE from tree structured documents.

iii) The use of negated literals in rule bodies improves significantly the precision of the IE rules in the case of using the *next* relation.

iv) The learnt IE rules were able to extract all the feature values (i.r. recall was 1) even when for training was used a single positive example which indicates a minimum user effort to manually extract a single item

from the training document.

v) The best result (with both precision and recall equal to 1) was obtained when we used for training all the fields to be extracted from the training document as positive examples and we allowed the presence of negated literals in rule bodies.

7. Related Work. With the rapid expansion of the Internet and the Web, the field of information extraction from HTML attracted a lot of researchers during the last decade. Clearly, it is impossible to mention all of their work here. However, at least two papers that are more closely related to the work reported in our paper, can be cited – (Freitag, 1998) and (Chidlovskii, 2003).

Paper (Freitag, 1998) describes the relational learner SRV that uses a FOIL-like algorithm for learning first order IE rules from a text document represented as a sequence of lexical tokens. The relations used in the rule bodies check various token features like: length, position in the text fragment, if they are numeric or capitalized, a.o. SRV has been adapted to learn IE rules from HTML. For this purpose new token features have been added to check the HTML context in which a token occurs. The most important similarity between SRV and our approach is the use of relational learning and a FOIL-like algorithm. The difference is that our approach has been explicitly devised to cope with tree structured documents, rather than string documents.

Paper (Chidlovskii, 2003) describes a generalization of the notion of string delimiters developed for IE from string documents ((Kushmerick, 2000)) to subtree delimiters for IE from tree documents. The paper describes a special purpose learner that constructs a structure called candidate index based on trie data structures, which is very different from FOIL’s approach. Note however, that the tree leaf delimiters described in that paper are very similar to our IE rules. Moreover, the representation of reverse paths using the symbols $Up(\uparrow)$, $Left(\leftarrow)$ and $Right(\rightarrow)$ can be easily simulated in our IE rules using the relations *child* and *next*.

8. Conclusions. This paper brings experimental evidence in support of the usefulness of first order rule induction for IE from HTML documents. As future research directions we would like to mention: i) the development of a GUI support for the IE approach outlined in this paper, in order to make the approach feasible for real users, ii) investigating the possibility of reducing the number of training examples required for learning good IE

rules and iii) experimenting with our approach in other application areas and for significantly larger collection of Web pages.

REFERENCES

- Bădică, C., A. Bădică, and V. Lițoiu (2003). *Enhancing WWW E-Commerce by Acquiring and Managing Product Knowledge*. In: *Proceedings of TAINN'03*, **E-7**, pp. 684-692, Çanakkale, Turkey.
- Chakrabarti, S. (2003). *Mining the Web. Discovering Knowledge from Hypertext Data*. Morgan Kaufmann Publishers, 344 pp.
- Chidlovskii, B. (2003). *Information Extraction from Tree Documents by Learning Subtree Delimiters*. In: *Proceedings of IJCAI-03 Workshop on Information Integration on the Web (IIWeb-03)*, pp. 3-8, Acapulco, Mexico.
- Fensel, D., Y. Ding, B. Omelayenko, E. Schulten, G. Botquin, M. Brown, A. Flet (2001). *Product Data Integration*. *IEEE Intelligent Systems*, **16(4)**, 54-59.
- Freitag, D. (1998). *Information extraction from HTML: application of a general machine learning approach*. In: *Proceedings of AAAI'98*, pp. 517-523.
- Kosala, R., J. van den Bussche, M. Bruynooghe, H. Blockeel. (2002). *Information Extraction in Structured Documents using Tree Automata Induction*. In: *Principles of Data Mining and Knowledge Discovery, 6th European Conference, LNAI 2431*, pp. 299-310.
- Kushmerick, N. (2000). Wrapper induction: Efficiency and expressiveness. *Artificial Intelligence*, **118**, 15-68.
- Lenhert, W., B. Sundheim (1991). A Performance Evaluation of Text-Analysis Technologies. In: *AI Magazine*, **12(3)**, 81-94.
- Mitchell, T.M. (1997). *Machine Learning*, McGraw-Hill, 414 pp.
- Muşlea, I. (1999). *Extraction Patterns for Information Extraction Tasks: A Survey*. In: *AAAI-99 Workshop on Machine Learning for Information Extraction*.
- Neven, F. (2002). Automata Theory for XML Researchers. *SIGMOD Record*, **31(3)**, 39-46.
- Quinlan, J.R., R.M. Cameron-Jones (1995). Induction of Logic Programs: FOIL and Related Systems, *New Generation Computing*, **13**, 287-312.

University of the West,
Faculty of Mathematics
Department of Computer Science
B-dul V. Pârvan, 4
Timișoara, 1900, ROMANIA
emailaddress@info.uvt.ro

Received January 2002