

Web Data Extraction, Applications and Techniques: A Survey

Emilio Ferrara^{a,*}, Pasquale De Meo^b, Giacomo Fiumara^c, Robert Baumgartner^d

^a*Center for Complex Networks and Systems Research, Indiana University, Bloomington, IN 47408, USA*

^b*Univ. of Messina, Dept. of Ancient and Modern Civilization, Polo Annunziata, I-98166 Messina, Italy*

^c*Univ. of Messina, Dept. of Mathematics and Informatics, viale F. Stagno D'Alcontres 31, I-98166 Messina, Italy*

^d*Lixto Software GmbH, Austria*

Abstract

Web Data Extraction is an important problem that has been studied by means of different scientific tools and in a broad range of applications. Many approaches to extracting data from the Web have been designed to solve specific problems and operate in ad-hoc domains. Other approaches, instead, heavily reuse techniques and algorithms developed in the field of Information Extraction.

This survey aims at providing a structured and comprehensive overview of the literature in the field of Web Data Extraction. We provided a simple classification framework in which existing Web Data Extraction applications are grouped into two main classes, namely applications at the Enterprise level and at the Social Web level. At the Enterprise level, Web Data Extraction techniques emerge as a key tool to perform data analysis in Business and Competitive Intelligence systems as well as for business process re-engineering. At the Social Web level, Web Data Extraction techniques allow to gather a large amount of structured data continuously generated and disseminated by Web 2.0, Social Media and Online Social Network users and this offers unprecedented opportunities to analyze human behavior at a very large scale. We discuss also the potential of cross-fertilization, i.e., on the possibility of re-using Web Data Extraction techniques originally designed to work in a given domain, in other domains.

Keywords:

Web Information Extraction, Web Data Mining, Business Intelligence, Knowledge Engineering, Knowledge-based Systems, Information Retrieval

*Corresponding author

Email addresses: ferrarae@indiana.edu (Emilio Ferrara), pdemeo@unime.it (Pasquale De Meo), gfiumara@unime.it (Giacomo Fiumara), robert.baumgartner@lixto.com (Robert Baumgartner)

Contents

1	Introduction	3
1.1	Challenges of Web Data Extraction techniques	3
1.2	Related work	4
1.3	Our contribution	4
1.4	Organization of the survey	5
2	Techniques	5
2.1	Tree-based techniques	6
2.1.1	Addressing elements in the document tree: XPath	6
2.1.2	Tree edit distance matching algorithms	7
2.2	Web wrappers	9
2.2.1	Wrapper generation and execution	10
2.2.2	The problem of wrapper maintenance	14
2.3	Hybrid systems: learning-based wrapper generation	16
3	Web Data Extraction Systems	17
3.1	The main phases associated with a Web Data Extraction System	17
3.2	Layer cake comparisons	19
4	Applications	21
4.1	Enterprise Applications	23
4.1.1	Context-aware advertising	23
4.1.2	Customer care	24
4.1.3	Database building	24
4.1.4	Software Engineering	24
4.1.5	Business Intelligence and Competitive Intelligence	24
4.1.6	Web process integration and channel management	25
4.1.7	Functional Web application testing	25
4.1.8	Comparison shopping	26
4.1.9	Mashup scenarios	26
4.1.10	Opinion mining	26
4.1.11	Citation databases	26
4.1.12	Web accessibility	27
4.1.13	Main content extraction	27
4.1.14	Web (experience) archiving	27
4.1.15	Summary	28
4.2	Social Web Applications	28
4.2.1	Extracting data from a single Online Social Web platform	30
4.2.2	Extracting data from multiple Online Social Web platforms	32
4.3	Opportunities for cross-fertilization	33
5	Conclusions	35

1. Introduction

Web Data Extraction systems are a broad class of software applications targeting at extracting information from Web sources [79, 11]. A Web Data Extraction system usually interacts with a Web source and extracts data stored in it: for instance, if the source is an HTML Web page, the extracted information could consist of elements in the page as well as the full-text of the page itself. Eventually, extracted data might be post-processed, converted in the most convenient structured format and stored for further usage [131, 63].

Web Data Extraction systems find extensive use in a wide range of applications including the analysis of text-based documents available to a company (like e-mails, support forums, technical and legal documentation, and so on), Business and Competitive Intelligence [9], crawling of Social Web platforms [17, 52], Bio-Informatics [99] and so on. The importance of Web Data Extraction systems depends on the fact that a large (and steadily growing) amount of information is continuously produced, shared and consumed online: Web Data Extraction systems allow to efficiently collect this information with limited human effort. The availability and analysis of collected data is an infeasible requirement to understand complex social, scientific and economic phenomena which generate the information itself. For example, collecting digital traces produced by users of Social Web platforms like Facebook, YouTube or Flickr is the key step to understand, model and predict human behavior [68, 94, 3].

In the commercial field, the Web provides a wealth of public domain information. A company can probe the Web to acquire and analyze information about the activity of its competitors. This process is known as *Competitive Intelligence* [22, 125] and it is crucial to quickly identify the opportunities provided by the market, to anticipate the decisions of the competitors as well as to learn from their faults and successes.

1.1. Challenges of Web Data Extraction techniques

The design and implementation of Web Data Extraction systems has been discussed from different perspectives and it leverages on scientific methods coming from various disciplines including Machine Learning, Logic and Natural Language Processing.

In the design of a Web Data Extraction system, many factors must be taken into account; some of them are independent of the specific application domain in which we plan to perform Web Data Extraction. Other factors, instead, heavily depend on the particular features of the application domain: as a consequence, some technological solutions which appear to be effective in some application contexts are not suitable in others.

In its most general formulation, the problem of extracting data from the Web is hard because it is constrained by several requirements. The key challenges we can encounter in the design of a Web Data Extraction system can be summarized as follows:

- Web Data Extraction techniques often require the help of human experts. A first challenge consists of providing a *high degree of automation* by reducing human efforts as much as possible. Human feedback, however, may play an important role in raising the level of accuracy achieved by a Web Data Extraction system.

A related challenge is, therefore, to identify a reasonable trade-off between the need of building highly automated Web Data Extraction procedures and the requirement of achieving accurate performance.

- Web Data Extraction techniques should be able to process large volumes of data in relatively short time. This requirement is particularly stringent in the field of Business and Competitive Intelligence because a company needs to perform timely analysis of market conditions.
- Applications in the field of Social Web or, more in general, those dealing with personal data must provide solid privacy guarantees. Therefore, potential (even if unintentional) attempts to violate user privacy should be timely and adequately identified and counteracted.

- Approaches relying on Machine Learning often require a significantly large training set of manually labeled Web pages. In general, the task of labeling pages is time-expensive and error-prone and, therefore, in many cases we can not assume the existence of labeled pages.
- Oftentimes, a Web Data Extraction tool has to routinely extract data from a Web Data source which can evolve over time. Web sources are continuously evolving and structural changes happen with no forewarning, thus are unpredictable. Eventually, in real-world scenarios it emerges the need of maintaining these systems, that might stop working correctly if lacking of flexibility to detect and face structural modifications of related Web sources.

1.2. Related work

The theme of Web Data Extraction is covered by a number of reviews. Laender et al. [79] presented a survey that offers a rigorous *taxonomy* to classify Web Data Extraction systems. The authors introduced a set of criteria and a qualitative analysis of various Web Data Extraction tools.

Kushmerick [77] defined a profile of *finite-state approaches* to the Web Data Extraction problem. The author analyzed both wrapper induction approaches (i.e., approaches capable of automatically generating wrappers by exploiting suitable examples) and maintenance ones (i.e., methods to update a wrapper each time the structure of the Web source changes). In that paper, Web Data Extraction techniques derived from Natural Language Processing and Hidden Markov Models were also discussed. On the wrapper induction problem, Flesca et al. [45] and Kaiser and Miksch [64] surveyed approaches, techniques and tools. The latter paper, in particular, provided a model describing the architecture of an Information Extraction system. Chang et al. [19] introduced a *tri-dimensional categorization* of Web Data Extraction systems, based on task difficulties, techniques used and degree of automation. In 2007, Fiumara [44] applied these criteria to classify four state-of-the-art Web Data Extraction systems. A relevant survey on Information Extraction is due to Sarawagi [105] and, in our opinion, anybody who intends to approach this discipline should read it. Recently, some authors focused on *unstructured data management systems (UDMSs)* [36], i.e., software systems that analyze raw text data, extract from them some structure (e.g. person name and location), integrate the structure (e.g., objects like *New York* and *NYC* are merged into a single object) and use the integrated structure to build a database. UDMSs are a relevant example of Web Data Extraction systems and the work from Doan et al. [36] provides an overview of *Cimple*, an UDMS developed at the University of Wisconsin. To the best of our knowledge, the survey from Baumgartner et al. [11] is the most recently updated review on the discipline as of this work.

1.3. Our contribution

The goal of this survey is to provide a structured and comprehensive overview of the research in Web Data Extraction as well as to provide an overview of most recent results in the literature.

We adopt a different point of view with respect to that used in other survey on this discipline: most of them present a list of tools, reporting a feature-based classification or an experimental comparison of these tools. Many of these papers are solid starting points in the study of this area. Unlike the existing surveys, our ambition is to provide a classification of existing Web Data Extraction techniques in terms of the application domains in which they have been employed. We want to shed light on the various research directions in this field as well as to understand to what extent techniques initially applied in a particular application domain have been later re-used in others. To the best of our knowledge, this is the first survey that deeply analyzes Web Data Extraction systems from a perspective of their application fields.

However, we also provide a detailed discussion of techniques to perform Web Data Extraction. We identify two main categories, i.e., approaches based on Tree Matching algorithms and approaches based on Machine Learning algorithms. For each category, we first describe the basic employed techniques and then we illustrate their variants. We also show how each category addresses the problems of wrapper generation and

maintenance. After that, we focus on applications that are strictly interconnected with Web Data Extraction tasks. We cover in particular enterprise, social and scientific applications by discussing which fields have already been approached (e.g., advertising engineering, enterprise solutions, Business and Competitive intelligence, etc.) and which are potentially going to be in the future (e.g., Bio-informatics, Web Harvesting, etc.).

We also discuss about the potential of *cross-fertilization*, i.e., whether strategies employed in a given domain can be re-used in others or, otherwise, if some applications can be adopted only in particular domains.

1.4. Organization of the survey

This survey is organized into two main parts. The first one is devoted to provide general definitions which are helpful to understand the material proposed in the survey. To this purpose, Section 2 illustrates the techniques exploited for collecting data from Web sources, and the algorithms that underlay most of Web Data Extraction systems. The main features of existing Web Data Extraction systems are largely discussed in Section 3.

The second part of this work is about the applications of Web Data Extraction systems to real-world scenarios. In Section 4 we identify two main domains in which Web Data Extraction techniques have been employed: applications at the enterprise level and at the Social Web level. The formers are described in Section 4.1, whereas the latters are covered in Section 4.2. This part concludes discussing the opportunities of cross-fertilization among different application scenarios (see Section 4.3).

In Section 5 we draw our conclusions and discuss potential applications of Web Data Extraction techniques that might arise in the future.

2. Techniques

The first part of this survey is devoted to the discussion of the techniques adopted in the field of the Web Data Extraction. In this part we extensively review approaches to extracting data from HTML pages. HTML is the predominant language for implementing Web pages and it is largely supported by W3C consortium. HTML pages can be regarded as a form of semi-structured data (even if less structured than other sources like XML documents) in which information follows a nested structure; HTML features can be profitably used in the design of suitable wrappers. However, we acknowledge that a large amount of semi-structured information is present in non-HTML formats (think of e-mail messages, software code and related documentations, system logs and so on) but the research approaches targeting at extracting information from this type of sources are out of the scope of this work.

The first attempts to extract data from the Web are dated back in early nineties. In the early days, this discipline borrowed approaches and techniques from Information Extraction (IE) literature. In particular, two classes of strategies emerged [64]: *learning techniques* and *knowledge engineering techniques* – also called *learning-based* and *rule-based* approaches, respectively [105]. These classes share a common rationale: the former was thought to develop systems that require human expertise to define rules (for example, *regular expressions*) to successfully accomplish the data extraction. These approaches require specific domain expertise: users that design and implement the rules and train the system must have programming experience and a good knowledge of the domain in which the data extraction system will operate; they will also have the ability to envisage potential usage scenarios and tasks assigned to the system. On the other hand, also some approaches of the latter class involve strong familiarity with both the requirements and the functions of the platform, so the human engagement is essential.

Various strategies have been devised to reduce the commitment of human domain experts. Some of them have been developed in the context of Artificial Intelligence literature, like the adoption of specific algorithms that use the structure of Web pages to identify and extract data. Others methods are borrowed from Machine

Learning, like supervised or semi-supervised learning techniques to design systems capable of being trained by examples and then able to autonomously extract data from similar (or even different) domains.

In the following we will discuss separately these two research lines. Section 2.1 presents the strategies based on the definition of algorithms capable of identifying information exploiting the semi-structured nature of Web pages. In Section 2.2 we introduce the concept of Web wrappers and explain how the techniques discussed in Section 2.1 are incorporated into Web wrappers; the then introduce some techniques for their generation and maintenance. Section 2.3 poses the attention on a new class of platforms that sit between approaches based on machine learning to induce wrappers and platforms to generate wrappers.

2.1. Tree-based techniques

One of the most exploited features in Web Data Extraction is the semi-structured nature of Web pages. These can be naturally represented as *labeled ordered rooted trees*, where labels represent the tags proper of the HTML mark-up language syntax, and the tree hierarchy represents the different levels of nesting of elements constituting the Web page. The representation of a Web page by using a labeled ordered rooted tree is usually referred as *DOM (Document Object Model)*, whose detailed explanation is out of the scope of this survey but has been largely regulated by the World Wide Web Consortium – For further details consider the following link: www.w3.org/DOM. The general idea behind the Document Object Model is that HTML Web pages are represented by means of plain text that contains HTML tags, particular keywords defined in the mark-up language, which can be interpreted by the browser to represent the elements specific of a Web page (e.g., hyper-links, buttons, images and so forth), so as free-text. HTML tags may be nested one into another, forming a hierarchical structure. This hierarchy is captured in the DOM by the document tree whose nodes represent HTML tags. The document tree (henceforth also referred as DOM tree) has been successfully exploited for Web Data Extraction purposes in a number of techniques discussed in the following.

2.1.1. Addressing elements in the document tree: XPath

One of the main advantage of the adoption of the Document Object Model for the HTML language is the possibility of exploiting some tools typical of XML languages (and HTML is to all effects a dialect of the XML). In particular, the XML Path Language (or, briefly, XPath) provides with a powerful syntax to address specific elements of an XML document (and, to the same extent, of HTML Web pages) in a simple manner. XPath has been defined by the World Wide Web Consortium, so as DOM – For the specifications see: www.w3.org/TR/xpath.

Although describing the syntax of XPath is not the core argument of this section, we provide Figure 1 as an example to explain how XPath can be used to address elements of a Web page. There exist two possible ways to use XPath: (i) to identify a single element in the document tree, or (ii) to address multiple occurrences of the same element. In the former case, illustrated in Figure 1(A), the defined XPath identifies just a single element on the Web page (namely, a table cell); in the latter, showed in Figure 1(B), the XPath identifies multiple instances of the same type of element (still a table cell) sharing the same hierarchical location.

To the purpose of Web Data Extraction, the possibility of exploiting such a powerful tool has been of utmost importance: the adoption of XPath as the tool to address elements in Web page has been largely exploited in the literature. The major weakness of XPath is related to its lack of flexibility: each XPath expression is strictly related to the structure of the Web page on top of which it has been defined. However, this limitation has been partially mitigated, by introducing *relative path expressions*, in latest releases (see XPath 2.0: www.w3.org/TR/xpath20). In general, even minor changes to the structure of a Web page might corrupt the correct functioning of an XPath expression defined on a previous version of the page.

To better clarify this concept, let us consider Web pages generated by a script (e.g., think of the information about a book in an e-commerce Web site). Now assume that the script undergoes some change: we can




Figure 1: Example of XPath(s) on the document tree, selecting one (A) or multiple (B) items.

expect that the tree structure of the HTML page generated by that script will change accordingly. To keep the Web Data Extraction process functional, one should update the expression every time any change occurs to the underlying page generation model; such an operation would require a high human commitment and, therefore, its cost could be prohibitively large. To this purpose, some authors [30, 31] introduced the concept of *wrapper robustness*: they proposed a strategy to find, among all the XPath expressions capable of extracting the same information from a Web page, the one that is less influenced by potential changes in the structure of the page and such an expression identifies the more robust wrapper. In general, to make the whole Web Data Extraction process robust, we need suitable tools allowing us to measure the similarity degree of two documents; such a task can be accomplished by detecting structural variations in the DOM trees associated with the documents. To this extent, some techniques, called *tree-matching strategies* are a good candidate to detect similarities between two trees and they will be discussed in detail in the next sections.

2.1.2. Tree edit distance matching algorithms

The first technique we describe is called *tree edit distance matching*. The problem of computing the tree edit distance between trees is a variation of the classic string edit distance problem. Given two labeled ordered rooted trees A and B , the problem is to finding a matching to transform A in B (or the vice-versa) with the minimum number of operations. The set of possible operations consists of node deletion, insertion or replacement. At each operation might be applied a cost, and in that case, the task turns in a cost-minimization problem (i.e., finding the sequence of operations of minimum cost to transform A in B).

The reasoning above is formally encoded in the definition of *mapping*, presented by [82]. A mapping M between two trees A and B is defined as a *set of ordered pairs* (i, j) , one from each tree, satisfying the following conditions $\forall (i_1, j_1), (i_2, j_2) \in M$.

1. $i_1 = i_2$ if and only if $j_1 = j_2$;
2. $A[i_1]$ is on the left of $A[i_2]$ if and only if $B[j_1]$ is on the left of $B[j_2]$;
3. $A[i_1]$ is an ancestor of $A[i_2]$ if and only if $B[j_1]$ is an ancestor of $B[j_2]$.

With the notation $A[i_x]$ we indicate the x -th node of the tree A in a *pre-order visit* of the tree. A number of consequences emerge from this definition of mapping:

- Each node must appear no more than once in a mapping;
- The order among siblings is preserved;
- The hierarchical relations among nodes is unchanged.

A number of techniques to approach this problem have been proposed [116, 23]. These methods support all three types of operations on nodes (i.e., node deletion, insertion and replacement) but are plagued by high computational costs. It has also been proved that the formulation for non ordered trees is NP-complete [128].

The simple tree matching algorithm. A computationally efficient solution for the problem of the tree edit distance matching is provided by the algorithm called *simple tree matching* [107], and its variants. This optimized strategy comes at a cost: node replacement is not allowed during the matching procedure – the shortcomings of this aspects will be further discussed below. The pseudo-code of the simple tree matching algorithm is provided in Algorithm 1, which adopts the following notation: $d(n)$ represents the degree of a node n (i.e., the number of first-level children); $T(i)$ is the i -th subtree of the tree rooted at node T .

Algorithm 1 SimpleTreeMatching(T' , T'')

```

1: if  $T'$  has the same label of  $T''$  then
2:    $m \leftarrow d(T')$ 
3:    $n \leftarrow d(T'')$ 
4:   for  $i = 0$  to  $m$  do
5:      $M[i][0] \leftarrow 0$ ;
6:     for  $j = 0$  to  $n$  do
7:        $M[0][j] \leftarrow 0$ ;
8:       for all  $i$  such that  $1 \leq i \leq m$  do
9:         for all  $j$  such that  $1 \leq j \leq n$  do
10:           $M[i][j] \leftarrow \text{Max}(M[i][j - 1], M[i - 1][j], M[i - 1][j - 1] + W[i][j])$  where  $W[i][j] =$ 
              SimpleTreeMatching( $T'(i - 1), T''(j - 1)$ )
11:   return  $M[m][n] + 1$ 
12: else
13:   return 0

```

The computational cost of simple tree matching is $O(\text{nodes}(A) \cdot \text{nodes}(B))$, where $\text{nodes}(T)$ is the function that returns the number of nodes in a tree (or a sub-tree) T ; the low cost ensures excellent performance when applied to HTML trees, which might be rich of nodes. Two main limitations hold with this algorithm:

- It can not match permutation of nodes;
- No level crossing is allowed (it is impossible to match nodes at different hierarchical levels).

Despite these intrinsic limits, this technique appears to fit very well to the purpose of matching HTML trees in the context of Web Data Extraction systems. In fact, it has been adopted in several scenarios [66, 126, 127, 130, 39, 40, 41]. One of the first Web Data Extraction approaches based on a tree edit distance algorithm is due to Reis *et al.* [102]. Such an approach focused on a very specific application domain (i.e, news extraction from the Web) but it was general enough to be re-applied in other domains. The algorithm of Reis *et al.* relies on a different definition of mapping called Restricted Top-Down Mapping (RTDM). In this model, insertion, removal and replacement operations are allowed only to the leaves of the trees. The restricted top-down edit distance between two trees A and B is defined as the cost of the restricted top-down mapping between the two trees. To find the restricted top-down mapping between two trees A and B , the Yang’s algorithm is applied [122]. The worst case time complexity of the approach of [102] is still $O(\text{nodes}(A) \cdot \text{nodes}(B))$ but, as shown by the authors, it works much faster because only restricted top-down mappings are managed. A normalized variant of the simple tree matching is called *normalized simple tree matching*. The normalization is computed as follows

$$NSTM(A, B) = \frac{\text{SimpleTreeMatching}(A, B)}{(\text{nodes}(A) + \text{nodes}(B))/2}.$$

The tree Matching algorithm and its variants are widely used in practice because they are easy to implement.

The weighted tree matching algorithm. Another variant of the simple tree matching is discussed in the following, and is called *weighted tree matching*. It adjusts the similarity values provided by the original

simple tree matching by introducing a re-normalization factor. The pseudo-codify of the weighted tree matching, recently presented in [39], is reported in Algorithm 2.

Algorithm 2 WeightedTreeMatching(T' , T'')

```

1: {Change line 11 in Algorithm 1 with the following code}
2: if  $m > 0$  AND  $n > 0$  then
3:   return  $M[m][n] * 1 / \text{Max}(t(T'), t(T''))$ 
4: else
5:   return  $M[m][n] + 1 / \text{Max}(t(T'), t(T''))$ 

```

In Algorithm 2, the notation $t(n)$ represents the number of total siblings of a node n including itself. Note that Algorithm 2 reports the differential version with respect to the simple tree matching described in Algorithm 1. The advantage of the weighted tree matching is that it better reflects a measure of similarity between two trees. In fact, in the simple tree matching algorithm the assigned matching value is always equal to one. Instead, the weighted tree matching algorithm assumes that less importance (i.e., a lower weight) is assigned to changes in the structure of the tree, when they occur in deeper sub-levels. This kind of changes can be, for example, missing or added leaves, truncated or added branches, etc. Also, a lower weight is accounted when changes occur in sub-levels with many nodes. The *weighted tree matching* algorithm returns a value in the interval $[0,1]$ and the closer to 1 the final value, the more similar the two input trees.

Let us analyze the behavior of the algorithm with an example often used in the literature [122, 126, 39] to explain the simple tree matching (see Figure 2). In that figure, A and B are two very simple generic rooted labeled trees (i.e., the same structure of HTML document trees). They show several similarities except for some missing nodes/branches. By applying the *weighted tree matching* algorithm, a value of $\frac{1}{3}$ is established for nodes (h), (i) and (j) belonging to A , although two of them are missing in B . Going up to parents, the summation of contributions of matching leaves is multiplied by the relative value of each node (e.g., in the first sub-level, the contribution of each node is $\frac{1}{4}$ because of the four first-sublevel nodes in A). Once completed these operations for all nodes of the sub-level, values are added and the final measure of similarity for the two trees is obtained.

In this example, the weighted tree matching between A and B returns a measure of similarity of $\frac{3}{8}$ (0.375) whereas the simple tree matching would return 7. The main difference on results provided by these two algorithms is the following: the *weighted tree matching* intrinsically produces a proper measure of similarity between the two compared trees while the *simple tree matching* returns the mapping value. A remarkable feature of the *weighted tree matching* algorithm is that, the more the structure of considered trees is complex and similar, the more the measure of similarity will be accurate.

2.2. Web wrappers

In the previous section we discussed some algorithms that might be adopted to identify information exploiting the semi-structured format of HTML documents. In the following we will discuss those procedures that might adopt the techniques presented above to carry out the data extraction.

In the literature, any procedure that aims at extracting structure data from unstructured (or semi-structured) data sources is usually referred as *wrapper*. In the context of Web Data Extraction we provide the following definition:

Definition 1 (Web wrapper). *A procedure, that might implement one or many different classes of algorithms, which seeks and finds data required by a human user, extracting them from unstructured (or semi-structured) Web sources, and transforming them into structured data, merging and unifying this information for further processing, in a semi-automatic or fully automatic way.*

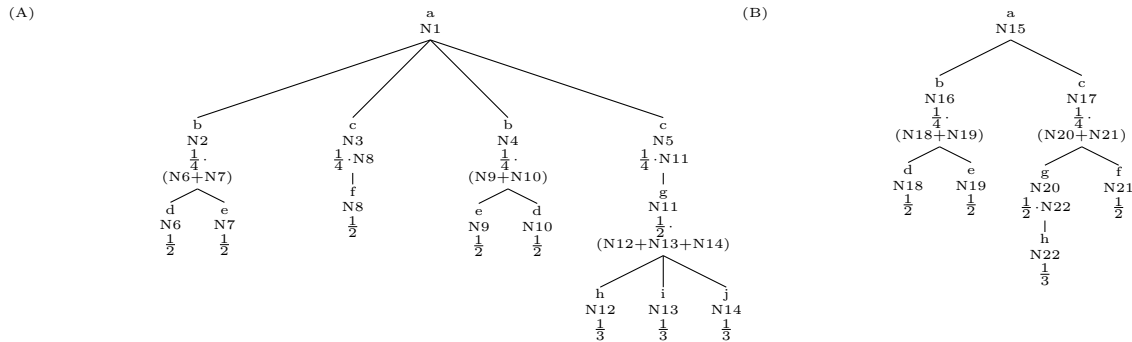


Figure 2: Example of application of the weighted tree matching algorithm for the comparison of two labeled rooted trees, A and B .

Web wrappers are characterized by a life-cycle constituted by several steps:

1. Wrapper generation: the wrapper is defined according to some technique(s);
2. Wrapper execution: the wrapper runs and extracts information continuously;
3. Wrapper maintenance: the structure of data sources may change and the wrapper should be adapted accordingly to keep working properly.

In the remainder of this section we will discuss these three different phases.

In particular, the first two steps of a wrapper life-cycle, its generation and execution, are discussed in Section 2.2.1; these steps might be implemented *manually*, for example by defining and executing regular expressions over the HTML documents; alternatively, which is the aim of Web Data Extraction systems, wrappers might be defined and executed by using an inductive approach – a process commonly known as *wrapper induction* [75]. Web wrapper induction is challenging because it requires high level automation strategies. There exist also hybrid approaches that make possible for users to generate and run wrappers semi-automatically by means of visual interfaces.

The last step of a wrapper life-cycle is the maintenance: Web pages change their structure continuously and without forewarning. This might corrupt the correct functioning of a Web wrapper, whose definition is usually tightly bound to the structure of the Web pages adopted for its generation. Defining automatic strategies for wrapper maintenance is of outstanding importance to guarantee correctness of extracted data and robustness of Web Data Extraction platforms. Some methodologies have been recently presented in the literature, as discussed in Section 2.2.2.

2.2.1. Wrapper generation and execution

The first step in wrappers life-cycle is their generation. Early Web Data Extraction platforms provided only support for manual generation of wrappers, which required human expertise and skills in programming languages to write scripts able to identify and extract selected pieces of information within a Web page.

In late nineties they made their appearance more advanced Web Data Extraction systems. The core feature provided was the possibility for their users to define and execute Web wrappers by means of interactive graphical users interfaces (GUIs). In most cases, it was not required any deep understanding of a wrapper programming language, as wrappers were generated automatically (or semi-automatically) by the system exploiting directives given by users by means of the platform interface.

In the following we discuss in detail three types of rationales underlying these kind of platforms, namely *regular expressions*, *wrapper programming languages* and *tree-based approaches*. The details regarding the features and of different Web Data Extraction platforms, instead, will be described in great detail in Section 3.

Regular-expression-based approach. One of the most common approaches is based on regular expressions, which are a powerful formal language used to identify strings or patterns in unstructured text on the basis of some matching criteria. Rules could be complex so, writing them manually, could require much time and a great expertise: wrappers based on regular expressions dynamically generate rules to extract desired data from Web pages. Usually, writing regular expressions on HTML pages relies on the following criteria: word boundaries, HTML tags, tables structure, etc. The advantage of platforms relying on regular expressions is that the user can usually select (for example by means of a graphical interface) one (or multiple) element(s) in a Web page, and the system is able to automatically infer the appropriate regular expression to identify that element in the page. Then, a wrapper might be created so that to extract similar elements, from other Web pages with the same structure of the one adopted to infer the regular expressions.

A notable tool implementing regular-expression-based extraction is W4F [104]. W4F adopts an annotation approach: instead of challenging users to deal with the HTML documents syntax, W4F eases the design of the wrapper by means of a *wizard procedure*. This wizard allows users to select and annotate elements directly on the Web page. W4F produces the regular expression extraction rules of the annotated items and provides them to users. A further step, which is the optimization of the regular expressions generated by W4F, is delegated to expert users – in fact, the tool is not always able to provide the best extraction rule. By fully exploiting the power of regular expressions, W4F extraction rules include *match* and also *split* expressions, which separates words, annotating different elements on the same string. The drawback of the adoption of regular expressions is their lack of flexibility. For example, whenever even a minor change occurs in the structure or content of a Web page, each regular expression is very likely to stop working, and must be rewritten. This process implies a big commitment by human users, in particular in the maintenance of systems based on regular expressions. For this reasons more flexible and powerful languages have been developed to empower the capabilities of Web Data Extraction platforms.

Logic-based approach. One example of powerful languages developed for data extraction purposes comes from the Web specific *wrapper programming languages*. Tools based on wrapper programming languages consider Web pages not as simply text strings but as semi-structured tree documents, whereas the DOM of the Web page represents its structure where nodes are elements characterized by both their properties and their content. The advantage of such an approach is that wrapper programming languages might be defined to fully exploit both the semi-structured nature of the Web pages and their contents – the former aspect lacks in regular-expression-based systems.

The first powerful wrapping language has been formalized by Gottlob and Koch [55]. The information extraction functions implemented by this wrapping language rely on monadic datalogs over trees [56]. The authors demonstrated that monadic datalogs over tree are equivalent to monadic second order logic (MSO), and hence very expressive. However, unlike MSO, a wrapper in monadic datalogs can be modeled nicely in a visual and interactive step-by-step manner. This makes this wrapping language suitable for being incorporated into visual tools, satisfying the condition that all its constructs can be implemented through *corresponding visual primitives*.

A bit of flavor on the functioning of the wrapping language is provided in the following. Starting from the unranked labeled tree representing the DOM of the Web page, the algorithm re-labels nodes, truncates the irrelevant ones, and finally returns a subset of original tree nodes, representing the selected data extracted. The first implementation of this wrapping language in a real-world scenarios is due to Baumgartner et al. [8, 7]. They developed the *Elog* wrapping language that implements almost all monadic datalog information extraction functions, with some minor restrictions. The *Elog* language is used as the core extraction method of the Lixto Visual Wrapper system. This platform provides a GUI to select, through visual specification, patterns in Web pages, in hierarchical order, highlighting elements of the document and specifying relationships among them. Information identified in this way could be too general, thus the system allows users to add some restricting conditions, for example before/after, not-before/not-after, internal and range conditions. Finally, selected data are translated into XML documents, by using pattern names as XML element names, obtaining structured data from semi-structured Web pages.

Tree-based approach [partial tree alignment]. The last technique discussed in this part relates to wrapper generation and is called *partial tree alignment*. It has been recently formalized by Zhai and Liu [126, 127] and the authors also developed a Web Data Extraction system based on it. This technique relies on the idea that information in Web documents usually are collected in contiguous regions of the page, called *data record regions*. The strategy of partial tree alignment consists in identifying and extracting these regions. In particular, the authors take inspiration from tree matching algorithms, by using the already discussed *tree edit distance* matching (see Section 2.1.2). The algorithm works in two steps:

1. Segmentation;
2. Partial tree alignment.

In the first phase, the Web page is split in segments, without extracting any data. This pre-processing phase is instrumental to the latter step. In fact, the system not only performs an analysis of the Web page document based on the DOM tree, but also relies on visual cues (like in the spatial reasoning technique, see Section 2.3), trying to identify gaps between data records. This step is useful also because helps the process of extracting structural information from the HTML document, in that situations when the HTML syntax is abused, for example by using tabular structure instead of CSS to arrange the graphical aspect of the page.

In the second step, the partial tree alignment algorithm is applied to data records earlier identified. Each data record is extracted from its DOM sub-tree position, constituting the root of a new single tree. This, because each data record could be contained in more than one non-contiguous sub-tree in the original DOM tree. The partial tree alignment approach implies the alignment of data fields with certainty, excluding those that can not be aligned, to ensure a high degree of precision. During this process no data items are involved, because partial tree alignment works only on tree tags matching, represented as the minimum cost, in terms of operations (i.e., node removal, node insertion, node replacement), to transform one node into another one. The drawback of this characteristic of the algorithm is that its recall performance (i.e., the ability of recovering all expected information) might decay in case of complex HTML document structures. In addition, also in the case of the partial tree alignment, the functioning of this strategy is strictly related with the structure of the Web page at the time of the definition of the alignment. This implies that the method is very sensitive even to small changes, that might compromise the functioning of the algorithm and the correct extraction of information. Even in this approach, the problem of the maintenance arises with outstanding importance.

Machine Learning approaches. Machine Learning techniques fit well to the purpose of extracting domain-specific information from Web sources, since they rely on training sessions during which a system acquires a domain expertise. Machine Learning approaches require a training step in which domain experts provide some manually labeled Web pages, acquired from different websites but also in the same website. Particular attention should be paid to providing examples of Web pages belonging to the same domain but exhibiting different structures. This, because, even in the same domain scenario, templates usually adopted to generate dynamic contents Web pages, differ, and the system should be capable of learning how to extract information in these contexts. Statistical Machine Learning systems were also developed, relying on conditional models [98] or adaptive search [114] as an alternative solution to human knowledge and interaction. In the following we shortly describe some Web Data Extraction approaches relying on Machine Learning algorithms [76, 16, 92, 109, 46, 60, 93].

One of the early approaches is *WIEN* [76]. WIEN was based on different *inductive learning techniques* and it was capable of automatically labeling training pages, representing *de facto* a hybrid system whose training process implied low human engagement. The flip side of the high automation of WIEN was the big number of limitations related to its inferential system: for example, the data extraction process was not capable of dealing with missing values – a case that occurs on a frequent base and posed serious limitations on the adaptability of WIEN to real-world scenarios.

Rapier (Robust Automated Production of Information Extraction Rules) [16, 92] is a system designed to learn rules for extracting information from documents and its main advantage is, perhaps, the capability of

learning these rules directly from documents without prior parsing or any post-processing. Extraction rules are relatively simple and they make use of limited syntactic and semantic information.

On one hand, Rapiere rules are flexible because they are not restricted to contain a fixed number of words but, on the other hand, it is hard to recognize what rules are actually useful to perform data extraction. To this purpose, a learning algorithm has been developed to find effective rules and this algorithm is based on Inductive Logic Programming.

WHISK was introduced in [109]. It relies on a supervised learning algorithm that generates rules for extracting information from text documents. *WHISK* is able to handle a wide range of text documents ranging from highly structured documents (like HTML pages) to free text. The extraction rules considered in *WHISK* can be regarded as a special type of regular expressions that have two components: the former specifies the context in which a phrase has to be considered relevant, the latter specifies the exact delimiters of the phrase to be extracted (i.e., the bounds of the text that has to be extracted). Depending of the structure of a document, *WHISK* generates rule that rely on exactly one of the two components cited above. In particular, in case of free text it uses context-based rules whereas in case of structured text it uses delimiters. In addition, for all those documents whose structure lies between structured document and free text *WHISK* is able to use a combination of context-based and delimiter-based rules. *WHISK* uses a supervised learning algorithm to induce novel rules from a set of hand-tagged instances. In order to keep human effort limited, *WHISK* interleaves the learning of new rules and the annotation of new instances. The learning/annotation process is, therefore, iterative and rather than presenting arbitrary instances, *WHISK* presents instances that are near to the examples that can be managed by the rules *WHISK* has learned so far.

SRV was proposed by Freytag [46]. SRV takes as input a set of tagged documents and extracts some features describing the tokens that can be extracted from a document. Features are classified into *simple* if they map a token onto a categorical value and *relational* if they map a token onto another token. SRV is also able to manage features encoding the structural aspect of a document (e.g., if a token is a verb). Extraction rules can be expressed on the basis of available features. To construct new rules, SRV uses a Naive Bayes classifier in conjunction with a relational learner.

SoftMealy [60] was the first wrapper induction system specifically designed to work in the Web Data Extraction context. Relying on *non-deterministic finite state automata* (also known as finite-state transducers (FST)), SoftMealy uses a bottom-up inductive learning approach to learn extraction rules. During the training session the system acquires training pages represented as an automaton on all the possible permutations of Web pages: states represent extracted data, while state transitions represent extraction rules. SoftMealy's main strength was its novel method of internal representation of the HTML documents. In detail, during a pre-processing step, each considered Web page was encoded into tokens (defined according to a set of inferential rules). Then, tokens were exploited to define separators, considered as invisible borderlines between two consecutive tokens. Finally, the FST was fed by *sequence of separators*, instead of raw HTML strings (as in WIEN), so that to match tokens with contextual rules (defined to characterize a set of individual separators) to determine the state transitions. The advantages of SoftMealy with respect to WIEN are worth noting: in fact, the system was able to deal with a number of exception, such as missing values/attributes, multiple attribute values, variant attribute permutations and also with typos.

The last learning-based system discussed in this part is called STALKER [93]. It was a supervised learning system for wrapper induction sharing some similarities with SoftMealy. The main difference between these two systems is the specification of relevant data: in STALKER, a set of tokens is manually positioned on the Web page, so that to identify information that the user intend to extract. This aspect ensures the capability of STALKER of handling with empty values, hierarchical structures and non ordered items. This system models a Web page content by means of hierarchical relations, represented by using a tree data structure called *embedded catalog tree* (EC tree). The root of the EC tree is populated the sequence of all tokens (whereas, STALKER considers as token any piece of text or HTML tag in the document). Each child node is a sub-sequence of tokens inherited by its parent node. This implies that each parent node is a

super-sequence of tokens of its children. The super-sequence is used, at each level of the hierarchy, to keep track of the content in the sub-levels of the EC tree. The extraction of elements of interest for the user is achieved by inferring a set of extraction rules on the EC tree itself – a typical example of extraction rule inferred by STALKER is the construct `SkipTo(T)`, a directive that indicates, during the extraction phase, to skip all tokens until the first occurrence of the token `T` is found. The inference of extraction rules exploits the concept of landmarks, sequences of consecutive tokens adopted to locate the beginning and the end of a given item to extract. STALKER is also able to define *wildcards*, classes of generic tokens that are inclusive of more specific tokens.

2.2.2. The problem of wrapper maintenance

Wrapper generation, regardless the adopted technique, is one aspect of the problem of data extraction from Web sources. On the other hand, *wrapper maintenance* is equally important, so that Web Data Extraction platform may reach high levels of robustness and reliability, hand in hand with the level of automation and low level of human engagement. In fact, differently from static documents, Web pages dynamically change and evolve, and their structure may change, sometimes with the consequence that previously defined wrappers are no longer able to successfully extract data.

In the light of these assumptions, one could argue that wrapper maintenance is a critical step of the Web Data Extraction process. Even though, this aspect has not acquired lot of attention in the literature (much less than the problem of wrapper generation), unless latest years. In the early stage, in fact, wrapper maintenance was performed manually: users that usually design Web wrappers, were updating or rewriting these wrappers every time the structure of a given Web page was modified. The manual maintenance approach was fitting well for small problems, but becomes unsuitable if the pool of Web pages largely increases. Since in the enterprise scenarios regular data extraction tasks might involve thousand (or even more) Web pages, dynamically generated and frequently updated, the manual wrapper maintenance is not anymore a feasible solution for real-world applications.

For these reasons, the problem of automatizing the wrapper maintenance has been faced in recent literature. For example, the first effort in the direction of automatic wrapper maintenance has been presented by Kushmerick [77], who defined for first the concept of *wrapper verification*. The task of wrapper verification arises as a required step during wrapper execution, in which a Web Data Extraction system assess if defined Web wrappers work properly or, alternatively, their functioning is corrupted due to modifications to the structure of underlying pages. Subsequently, the author discussed some techniques of semi-automatic wrapper maintenance, to handle simple problems.

The first method that tries to automatize the process of wrapper maintenance has been developed by Meng et al. [89], and it is called *schema-guided wrapper maintenance*. It relies on the definition of XML schemes during the phase of the wrapper generation, to be exploited for the maintenance during the execution step. More recently, Ferrara and Baumgartner [39, 40, 41] developed a system of automatic wrapper adaptation (a kind of maintenance that occur to modify Web wrappers according to the new structure of the Web pages) relying on the analysis of structural similarities between different versions of the same Web page using a tree-edit algorithm. In the following, we discuss these two strategies of wrapper maintenance.

Schema-guided wrapper maintenance. The first attempt to deal with the problem of wrapper maintenance providing a high level of automation has been presented by Meng et al. [89]. The authors developed SG-WRAM (Schema-Guided WRApper Maintenance), a strategy for Web Data Extraction that is built on top of the assumption, based on empirical observations, that changes in Web pages, even substantial, often preserve:

- Syntactic features: syntactic characteristics of data items like data patterns, string lengths, etc., are mostly preserved.

- Hyperlinks: HTML documents are often enriched by hyperlinks that are seldom removed in subsequent modifications of the Web pages.
- Annotations: descriptive information representing the semantic meaning of a piece of information in its context is usually maintained.

In the light of these assumptions, the authors developed a Web Data Extraction system that, during the phase of wrapper generation, creates schemes which will be exploited during the phase of wrapper maintenance. In detail, during the generation of wrappers, the user provides HTML documents and XML schemes, specifying a mapping between them. Later, the system will generate extraction rules and then it will execute the wrappers to extract data, building an XML document according to the specified XML schema. During the wrapper execution phase, an additional component is introduced in the pipeline of the data extraction process: the *wrapper maintainer*. The wrapper maintainer checks for any potential extraction issue and provides an automatic repairing protocol for wrappers which fail their extraction task because of modifications in the structure of related Web pages. The repairing protocol might be successful, and in that case the data extraction continues, or it might fail – in that case warnings and notifications rise. The XML schemes are defined in the format of a DTD (Document Type Definition) and the HTML documents are represented as DOM trees, according to what explained in Section 2.1. The SG-WRAM system builds corresponding mappings between them and generates extraction rules in the format of XQuery expressions – For XQuery specifications see: www.w3.org/TR/xquery.

Automatic wrapper adaptation. Another strategy for the automatic maintenance of Web wrappers has been recently presented in [39, 40, 41]. In detail, it is a method of automatic wrapper adaptation that relies on the idea of comparing helpful structural information stored in the Web wrapper defined on the original version of the Web page, searching for similarities in the new version of the page, after any structural modification occurs.

The strategy works for different data extraction techniques implemented by the wrapping system. For example, it has been tested by using both XPath (see Section 2.1.1) and the Elog wrapping language (see Section 2.2.1). In this strategy, elements are identified and represented as sub-trees of the DOM tree of the Web page, and can be exploited to find similarities between two different versions of the same document. We discuss an example adopting XPath to address a single element in a Web page, as reported in the example in Figure 1(A).

The rationale behind the automatic wrapper adaptation is to search for some elements, in the modified version of the Web page, that share structural similarities with the original one. The evaluation of the similarity is done on the basis of comparable features (e.g., subtrees, attributes, etc.). These elements are called *candidates*: among them, the one showing the higher degree of similarity with the element in the original page, is matched in the new version of the page. The algorithm adopted to compute the matching among the DOM trees of the two HTML documents is the weighted tree matching, already presented in Section 2.1.2. Further heuristics are adopted to assess the similarity of nodes, for example exploiting additional features exhibited by DOM elements. In some cases, for example, their textual contents might be compared, according to string distance measures such as Jaro-Winkler [120] or bigrams [24], to take into account also the content similarity of two given nodes.

It is possible to extend the same approach in the case in which the XPath identifies multiple similar elements on the original page (e.g., an XPath selecting results of a search in a retail online shop, represented as table rows, divs or list items), as reported in Figure 1(B). In detail, it is possible to identify multiple elements sharing a similar structure in the new page, within a custom level of accuracy (e.g., establishing a threshold value of similarity).

The authors implemented this approach in a commercial tool – Lixto, describing how the pipeline of the wrapper generation has been modified to allow the wrappers to automatically detect and adapt their functioning to structural modifications occurring to Web pages [41]. In this context, the strategy proposed was

to acquire structural information about those elements the original wrapper extracts, storing them directly inside the wrapper itself. This is done, for example, generating *signatures* representing the DOM sub-tree of extracted elements from the original Web page, stored as a tree diagram, or as a simple XML documents. During the execution of Web wrappers, if any extraction issue occurs due to a structural modification of the page, the wrapper adaptation algorithm automatically starts and tries to adapt the wrapper to the new structure.

2.3. Hybrid systems: learning-based wrapper generation

Wrapper generation systems and wrapper induction techniques discussed above differ essentially in two aspects:

1. The degree of automation of the Web Data Extraction systems;
2. The amount and the type of human engagement required during operation.

The first point is related to the ability of the system to work in an autonomous way, ensuring sufficient standards of robustness and reliability, according to the requirements of users. Regarding the second point, most of the wrapper induction systems need labeled examples provided during the training sessions, thus requiring human expert engagement for the manual labeling phase. Wrapper generation systems, on the other hand, engage users into their maintenance, unless automatic techniques are employed, as those discussed in Section 2.2.2.

Interestingly, a new class of platforms has been discussed in recent literature, that adopts a hybrid approach that sits between learning-based wrapper induction systems and wrapper generation platforms. The first example of this class of systems is given by RoadRunner [28, 27], a template-based system that automatically generates templates to extract data by matching features from different pages in the same domain. Another interesting approach is that of exploiting visual cues and spatial reasoning to identify elements in the Web pages with a Computer Vision oriented paradigm. This part concludes with the discussion of these two systems.

Template-based matching. The first example of hybrid system is provided by RoadRunner [28, 27]. This system might be considered as an interesting example of automatic wrapper generator. The main strength of RoadRunner is that it is oriented to data-intensive websites based on templates or regular structures. The system tackles the problem of data extraction exploiting both features used by wrapper generators, and by wrapper induction systems. In particular, RoadRunner can work using information provided by users, in the form of labeled example pages, or also by automatically labeling Web pages (such as WIEN), to build a training set. In addition, it might exploits *a priori* knowledge on the schema of the Web pages, for example taking into account previously learned page templates. RoadRunner relies on the idea of working with two HTML pages at the same time in order to discover patterns analyzing similarities and differences between structure and content of each pair of pages. Essentially, RoadRunner can extract relevant information from any Web site containing at least two Web pages with a similar structure. Since usually Web pages are dynamically generated starting from template, and relevant data are positioned in the same (or in similar) areas of the page, RoadRunner is able to exploit this characteristic to identify relevant pieces of information, and, at the same time, taking into account small differences due to missing values or other mismatches.

The authors defined as *class of pages* those Web sources characterized by a common generation script. Then, the problem is reduced to extracting relevant data by generating wrappers for class of pages, starting from the inference of a common structure from the two-page-based comparison. This system can handle missing and optional values and also structural differences, adapting very well to all kinds of data-intensive Web sources. Another strength of RoadRunner is its high-quality open-source implementation (see: www.dia.uniroma3.it/db/roadRunner/), that provides a high degree of reliability of the extraction system.

Spatial reasoning. The paradigm of the Computer Vision has also inspired the field of Web Data Extraction systems. In fact, a recent model of data extraction, called Visual Box Model, has been presented [72, 51]. The Visual Box Model exploits visual cues to understand if, in the version of the Web page displayed on the screen, after the rendering of the Web browser, are present, for example, data in a tabular format. The advantage of this strategy is that it is possible to acquire data not necessarily represented by means of the standard HTML `<table>` format.

The functioning of this technique is based on a X-Y cut OCR algorithm. This algorithm is able, given the rendered version of a Web page, of generating a visual grid, where elements of the page are allocated according to their coordinates – determined by visual cues. Cuts are recursively applied to the bitmap image representing the rendering of the Web page, and stored into an X-Y tree. This tree is built so that ancestor nodes with leaves represent not-empty tables. Some additional operations check whether extracted tables contain useful information. This is done because – although it is a deprecated practice – many Web pages use tables for structural and graphical purposes, instead of for data representation scopes.

The Visual Box Model data extraction system is implemented by means of an internal rendering engine that produces a visualization of the Web page relying on Gecko (developer.mozilla.org/en/Gecko), the same rendering engine used by the Mozilla Web browser. By exploiting the CSS 2.0 box model, the algorithm is able to access the positional information of any given element. This is achieved by a bridge between the rendering engine and the application, implemented by means of the XPCOM library (developer.mozilla.org/it/XPCOM).

3. Web Data Extraction Systems

In this section we get into details regarding the characteristics of existing Web Data Extraction systems. We can generically define a Web Data Extraction system as a platform implementing *a sequence of procedures (for example, Web wrappers) that extract information from Web sources* [79]. A large number of Web Data Extraction systems are available as commercial products even if an increasing number of free, open-source alternatives to commercial software is now entering into the market. In most cases, the average end users of Web Data Extraction systems are firms or data analysts looking for relevant information from the Web. An intermediate category of users consists of non-expert individuals that need to collect some Web content, often on a non-regular basis. This category of users are often non-experienced and they look at simple but powerful Web Data Extraction software suites; among them we can cite DEiXTo¹; DEiXTo is based on the W3C Document Object Model and enable users to easily create extraction rules pointing out the portion of data to scrape from a Web site.

In the next subsections we first illustrate the various phases characterizing a Web Data Extraction system (see Section 3.1) Subsequently, in Section 3.2, we consider various factors influencing the design and implementation of Web Data Extraction system (e.g. the task of generating wrappers according to the easiness of use). We then illustrate the technological evolution of Web Data Extraction systems according to each factor under inquiry. To this purpose, we use ad hoc diagrams consisting of a group of layers (*layer cakes*) such that bottom (resp., top) layers correspond to the earliest (resp., latest) technological solutions.

3.1. The main phases associated with a Web Data Extraction System

In this section we describe the various phases associated with the procedure of extracting data from a Web site.

¹<http://deixto.com>

Interaction with Web pages. The first phase of a generic Web Data Extraction system is the *Web interaction* [117]: the Web Data Extraction system accesses a Web source and extracts information stored in it. Web sources usually coincide with Web pages, but some approaches consider also as RSS/Atom feeds [57] and Microformats [65].

Some commercial systems, Lixto for first but also Kapow Mashup Server (described below), include a *Graphical User Interface* for fully *visual and interactive navigation* of HTML pages, integrated with data extraction tools.

The most advanced Web Data Extraction systems support the extraction of data from pages reached by *deep Web navigation* [6]: they simulate the activity of users clicking on DOM elements of pages, through macros or, more simply, by filling HTML forms.

These systems also support the extraction of information from *dynamically generated Web pages*, usually built at run-time as a consequence of the user request, filling a template page with data from some database. The other kind of pages are commonly called *static Web pages*, because of their static content.

OXPath [48], which is part of the DIADEM project [47] is a declarative formalism that extends XPath to support deep Web navigation and data extraction from interactive Web sites. It adds new location steps to simulate user actions, a new axis to select dynamically computed CSS attributes, a convenient field function to identify visible fields only, page iteration with a Kleene-star extension for repeated navigations, and new predicates for marking expressions for extraction identification.

Generation of a wrapper. A Web Data Extraction system must implement the support for wrapper generation and wrapper execution.

Another definition of Web Data Extraction system was provided by Baumgartner et al. [11]. They defined a Web Data Extraction system as “a software extracting, *automatically* and *repeatedly*, data from Web pages with changing contents, and that *delivers extracted data* to a database or some other application”. This is the definition that better fits the modern view of the problem of the Web Data Extraction as it introduces three important aspects:

- Automation and scheduling
- Data transformation, and the
- Use of the extracted data

In the following we shall discuss each of these aspects into detail.

Automation and Extraction. Automating the access to Web pages as well as the localization of their elements is one of the most important features included in last Web Data Extraction systems [98].

Among the most important automation features we cite the possibility to *simulate the click stream of the user*, *filling forms* and *selecting menus and buttons*, the support for AJAX technology [49] to handle the asynchronous updating of the page and the ability of scheduling Web data extraction procedures on a periodical basis.

Data transformation. Information could be wrapped from multiple sources, which means using different wrappers and also, probably, obtaining different structures of extracted data. The steps between extraction and delivering are called *data transformation*: during these phases, such as *data cleaning* [101] and *conflict resolution* [91], users reach the target to obtain homogeneous information under a unique resulting structure. The most powerful Web Data Extraction systems provide tools to perform *automatic schema matching* from multiple wrappers [100], then packaging data into a desired format (e.g., a database, XML, etc.) to make it possible to *query data*, *normalize structure* and *de-duplicate tuples*.

Use of extracted data. When the extraction task is complete, and acquired data are packaged in the needed format, this information is ready to be used; the last step is to deliver the package, now represented by

structured data, to a managing system (e.g., a native XML DBMS, a RDBMS, a data warehouse, a CMS, etc.). In addition to all the specific fields of application covered later in this work, acquired data can be also generically used for analytical or statistical purposes [14] or simply to republish them under a structured format.

3.2. Layer cake comparisons

In this section, we summarize the capability stacks of Web data extraction systems from our understanding, including aspects of wrapper generation, data extraction capabilities, and wrapper usage. In particular, we introduce some specific aspects and illustrate the technological evolution of Web Data Extraction systems according to each aspect. We use ad hoc diagrams structured as a group of layers (*layer cakes*). In these diagrams, bottom (resp., top) layers correspond to the earliest (resp., latest) technological solutions.

Wrapper Generation: Ease of Use (Figure 3). The first approaches to consuming facts from the Web were implemented by means of general purpose languages. Over time libraries (e.g. Ruby Mechanize) and special-purpose query languages evolved on top of this principle (e.g., Jedi [61] and Florid [87]). Wizards that simplify the way to specify queries are the next logical level and for instance have been used in W4F [104] and XWrap [83]. Advanced Web Data Extraction systems offer GUIs for configuration, either client-based (e.g. Lapis), Web-based (e.g. Dapper and Needlebase) or as browser extensions (e.g. iOpus and Chickenfoot). Commercial frameworks offer a full IDE (Integrated Development Environment) with functionalities described in the previous section (e.g. Denodo, Kapowtech, Lixto and Mozenda).

Wrapper Generation: Creation Paradigma (Figure 4). From the perspective of how the system supports the wrapper designer to create robust extraction programs, the simplest approach is to manually specify queries and test them against sample sites individually. Advanced editors offer highlighting of query keywords and operators and assist query writing with auto-completion and similar usability enhancements (e.g., Screen-Scraper). In case of procedural languages, debugging means and visual assistance for constructs such as loops are further means to guided wrapper generation. Regarding the definition of deep Web navigations, such as form fillouts, a number of tools offer VCR-style recording the human browsing and replaying the recorded navigation sequence (e.g., Chickenfoot, iOpus, Lixto). Visual and interactive facilities are offered by systems to simplify fact extraction. Users mark an example instance, and the system identifies the selected element in a robust way, and possibly generalize to match further similar elements (e.g., to find all book titles). Such means are often equipped with Machine Learning methods, in which the user can select a multitude of positive and negative examples, and the system generates a grammar to identify the Web objects under consideration, often in an iterative fashion (e.g. Wien, Dapper, Needlebase). Click and drag and drop features further simplify the interactive and visual mechanisms. Finally, some systems offer vertical templates to easily create wrappers for particular domains, for example for extracting hotel data or news item, using Natural Language Processing techniques and domain knowledge, or for extracting data from typical Web layouts, such as table structures or overview pages with next links.




Figure 3: Wrapper Generation: Ease of Use




Figure 4: Wrapper Creation Paradigma

Deep Web Navigation Capabilities (Figure 5). Before the advent of Web 2.0 techniques, dynamic HTML and AJAX it was usually sufficient to consider the Web as a collection of linked pages. In such cases, form filling can be simulated by tracking the requests and responses from the Web Server, and replaying the sequence of requests (sometimes populating a session id dynamically, extracted from a previous page of the sequence). Alternatively, early Web Data extraction systems have been influenced by screen scraping technologies like they were being used for automating 3270 applications or like used for automating native applications, usually relying heavily on coordinates. Understanding and replaying DOM Events on Web objects is the next logical level in this capability stack. Advanced systems even go a step further, especially when embedding a full browser: the click on an element is recorded in a robust way and during replay the browser is informed to do a visual click on such an element, handing the DOM handling over to the browser and making sure that the Web page is consumed exactly in the way the human user consumes it. Orthogonal to such features are capabilities to parametrize deep Web sequences, and to use query probing techniques to automate deep Web navigation to unknown forms.

Web Data Extraction Capabilities (Figure 6). Over time, various approaches to modeling a Web page have been discussed. The simplest way is to work on the stream received from the Web server, for example using regular expressions. In some cases, this is sufficient and even the preferred approach in large-scale scenarios due to avoiding to build up a complex and unperformant model. On the other hand, in complex Web 2.0 pages or in pages that are not well-formed, it can be extremely cumbersome to work on the textual level only. Moreover, such wrappers are not very maintainable and break frequently. The most common approach is to work on the DOM tree or some other kind of tree structure. This approach has been followed both by the academic community and by commercial approaches. In the academic communities studying expressiveness of language over trees and tree automata are interesting, and from a commercial point of view, it is convenient and robust to use languages such as XPath for identifying Web objects. Usually, not only the elements of a DOM tree are considered, but also the events, enabling to specify data extraction and navigation steps with the same approach. In the Web of today, however, very often the DOM tree does not really capture the essential structure of a Web page as presented to the human user in a browser. A human perceives something as table structure, whereas the DOM tree contains a list of `div` elements with absolute positioning. Moreover, binary objects embedded into Web pages such as Flash pose further challenges and are not covered with a tree structure approach. Hence, screen-scraping made back its way into novel Web Data extraction frameworks, using methods from document understanding and spatial reasoning such as the approaches of the TamCrow project [73], of the ABBA project [37] spatial XPath extensions [95] and rendition-based extensions in RoadRunner to detect labels [29].

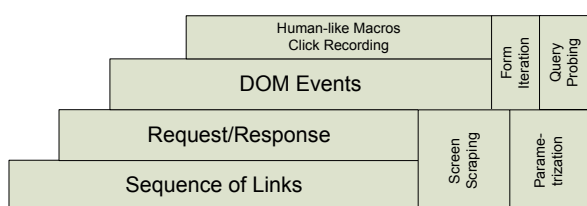


Figure 5: Deep Web Navigation Capabilities




Figure 6: Web Data Extraction Capabilities

Parser and Browser Embedding (Figure 7). This capability stack is closely related to the Deep Web capabilities but focussing on the technical realization of parsing and browser embedding. Simple approaches create their own parser to identify relevant HTML tags, more sophisticated approaches use DOM libraries without an associated browser view. Due to the fact that many Web Data Extraction frameworks are implemented in Java, special-purpose browsers such as the Java Swing browser and the ICE browser are and have been used. The most powerful approaches are the ones that embed a standard browser such as IE, Firefox or WebKit-based browsers. In case of Java implementations, interfaces such as Java-XPCOM bridge or libraries such as JREx are used to embed the Mozilla browser. Embedding a full browser not only gives

access to the DOM model, but additionally to other models useful for data extraction, including the CSS Box model. Some tools go a different direction and instead of embedding a browser, they are implemented as browser extension, imposing some restrictions and inconveniences. Orthogonal to the browser stack are the capabilities to extend extraction functionalities to unstructured text parsing exploiting Natural Language Processing techniques.

Complexity of Supported Operations (Figure 8). Simple data extraction tools offer Web notification, e.g. if a particular word is mentioned. Web macro recorders allow users to create “deep” bookmarks, and Web clipping frameworks clip fragments of a Web page to the desktop. Personalization of a Web page (e.g., some tools offer to alter the CSS styles to make frequently used links more prominent on a page) is a further level in the layer cake. Batch processing frameworks offer functionalities to replay ample extraction tasks (e.g., running through many different values in form fillout). Advanced systems go a step further and use sophisticated extraction plans and anonymization techniques that ensure to stay under the radar and not harm target portals with too many requests at once, and deliver the data into further applications such as market intelligence platforms. The expressiveness of a wrapper language contributes as well to the complexity of supported operations.




Figure 7: Parser and Browser Embedding




Figure 8: Complexity of Supported Operations

4. Applications

The aim of the second part of this paper is to survey and analyze a large number of applications that are strictly interconnected with Web Data Extraction tasks. To the best of our knowledge, this is the first attempt to classify applications based on Web Data Extraction techniques even if they have been originally designed to operate in specific domain and, in some cases, they can appear as unrelated.

The spectrum of applications possibly benefiting from Web Data Extraction techniques is quite large. We decided to focus on two main application scenarios, namely the *enterprise domain* and the *Social Web domain* because, in our opinion, these scenarios can take advantage the most from Web Data Extraction technologies. As for the enterprise domain, Web Data Extraction software can significantly increase the efficiency of business processes by freeing up resources like time and manpower. Due to the automatic collection of data from the Web, companies have at their disposal large amount of data in a relatively short time frame and these data assist firm analysts and manager to plan or revise business strategies. As for the Social Web domain, extracting data from Social Web sites/Social Media is a relatively new concept that has acquired a big relevance due to the explosive growth of platforms like Facebook, Twitter or Instagram. Crawling social media and online social networks data has become the flagship Social application of Web Data Extraction [17, 42]. This yield the unprecedented opportunity for scientists from different domains (e.g., Sociology, Political Science, Anthropology, etc.) to analyze and understand the dynamics of human behaviors at a planetary scale and in a real time-fashion [38, 25, 33, 43]. The classification scheme above, however, has not to be intended as hard because there are some applications which have been originally designed to work in the enterprise domain but they have subsequently (and successfully) re-used in the context of the Social Web. For instance, companies can extract data from Social Web platforms to pursuit tasks like *reputation management* (i.e., a company is informed of who is talking about the firm itself and its

products) and *trend management* (i.e., a company can identify what are the topics catching the interest of its customers). In general, user generated content available on Social Web platforms are perhaps the data that firms are interested the most: data such as product or service reviews, feedback and complaints are a key ingredient to monitor the impact of a brand on the market and the sentiments toward that brand and to detect how a brand is perceived in opposition to its commercial competitors. On the other hand, data produced by users who discuss on the same product/service/brand have also a social dimension because they reflect social interactions: a user in a given community (e.g., a discussion forum of a company) can pose questions to other members on the same community and they can rank the advice they receive on the basis of the social relationship like trust or friendship binding the user who pose the question with the users who answered it.

In addition to classifying applications into the enterprise and Social Web categories, we can introduce other criteria which rely on the fact that data are extracted from a single or multiple sources, data can be of the same format or, finally, are associated with the same domain. In particular, we introduce the following criteria:

Single Source vs. Multiple Sources. Web Data Extraction techniques can be applied on data residing on a single platform or, vice versa, they can collect data located in different platforms.

As an example, in the context of Enterprise applications, some applications fetch data from a single platform: a relevant example is provided by applications to manage customer care activities. As previously pointed out, these application crawls a corpus of text documents produced within a single platform. By contrast, in some cases, the application can benefit of data scattered across multiple system. In the context of Enterprise applications, a notable example of application gathering data from multiple sources is given by *Comparison Shopping*: in such a case, a Web Data Extraction technique is able to collect data associated with a target product from a multitude of e-commerce platforms and to compare, for instance, the price and the shipping details of that product in each of these platform.

The classification above is particularly interesting in the field of Social Web applications: in fact we discuss separately applications designed to crawl and collect data from a single platform (see Section 4.2.1) from applications running on multiple Social Web platforms.

Homogeneity in Format vs. Heterogeneity in Format. The second classification criterium we introduce answers the following question: does the application collect data of the same format or, vice versa, can it collect data of different formats? On the basis of such a criterium we can classify existing applications in *homogeneous-format* and *heterogeneous-format*.

As for enterprise applications, a nice example of homogeneous-format applications is given by *Opinion Sharing Applications*. This term identifies applications devoted to collect the opinions a user expresses about a given service/product. These opinions are typically represented by short textual reviews or blog post. In such a case, the type of data extracted by a Web Data Extraction technique is a string. In some cases, users are allowed to provide scores and the format of extracted data is a discrete value, generally ranging in an interval like $[0,5]$. In case of Social Web applications, many approaches (like [52, 17, 90, 123]) are devoted to extract friendship relationships from a Social Networking Websites. In such a case, extracted data are of the same format that can be interpreted as a triplet of the form $\langle u_x, u_y, f \rangle$ being u_x and u_y the user identifiers of two Social Network members; the value f can be a boolean (e.g., it can be `true` if u_x and u_y are friends, `false` otherwise) or a numerical value (e.g., it can be set equal to the number of messages u_x sent u_y).

Some applications are designed to collect data of different type. A relevant example in the field of enterprise applications is provided by those applications designed to manage Business Intelligence tasks. These application are able to collect data of different type like (numerical data or textual ones) as well as to manage both structured data (e.g., data coming from relational tables of a DBMS) as well as unstructured ones (e.g., text snippets present in HTML pages). Relevant examples of applications capable of extracting data

of different type are also present in the context of Social Web applications: for instance, [78] performed in 2009 a crawl of the whole Twitter platform which produced textual data (e.g., the tweets and re-tweets produced by users) as well as data indicating the different types of connections among users (“following”, “reply to” and “mention”) [103].

Single Purpose or Multiple Purpose. In some cases, the goal of a Web Data Extraction tool is to extract data describing a specific facet of a particular social phenomenon or business process. In such a case, an application is catalogued as *unique purpose*. Other applications, instead, aim at collecting data of different nature which, if properly linked, are relevant to better understand and interpret a particular phenomenon. After linking data, novel and sophisticated applications running on them can be defined. Applications belonging to this category will be called *multi purpose*. As an example of single purpose applications, we focus on applications devoted to collect bibliographic data as well as citations among papers. To collect bibliographical information, Web Data Extraction techniques are required to query several databases containing scientific publications (like ISI Web of Science, SCOPUS, PubMed and so on) but the goal of the application, however, is to collect all the citations associated with a particular paper or to a particular author because this information is useful to assess the impact of an author or a scientific publication in a scientific discipline. In the field of Social Web applications, an example of single purpose application is given by applications devoted at collecting data about human activities in different Social Web platforms: for instance, the tags contributed by a user in different systems [112]. This is relevant to understand if the language of a user is uniform across various platforms or if the platform features impact on the user formation of the vocabulary.

We can cite some relevant examples of multi purpose applications both at the enterprise and at the Social Web level. For instance, in the class of enterprise applications, some applications are able to collect data produced by different Web Services and combine these data to produce more advanced applications. For instance, in the travel industry, one can think of applications collecting data on flights and hotels and combine these data to produce holiday packages. In the field of Social Web applications, several authors observed that the merging of different type of data describing different type of human activities produces a more detailed knowledge of human needs and preferences. For instance, in [106] the authors conducted extensive experiments on data samples extracted from Last.Fm and Flickr. They showed that a strong correlations exist between user social activities (e.g., the number of friends of a user or the number of groups she joined) and the tagging activity of the same user. A nice result provided in [106] was that user contributed tags are a useful indicator to predict friendship relationships. Other studies entwining users geographical locations with the content they posted are provided in [26] and [67].

4.1. Enterprise Applications

In this section we describe the main features of software applications and procedures related with Web Data Extraction with a direct or subsequent *commercial scope*.

4.1.1. Context-aware advertising

Context-aware advertising techniques aim at presenting to the end user of a Web site commercial *thematized advertisements* together with the content of the Web page the user is reading. The ultimate goal is to increase the value that a Web page can have for its visitors and, ultimately, to raise the level of interest in the ad.

First efforts to implement context-aware advertising were made by Applied Semantic, Inc. (www.appliedsemantics.com); subsequently Google bought their “AdSense” advertising solution. The implementation of context-aware advertisements requires to *analyze the semantic content* of the page, *extract relevant information*, both in the structure and in the data, and then contextualize the ads content and placement in the same page. Contextual advertising, compared to the old concept of Web advertising, represents a *intelligent approach* to

providing useful information to the user, statistically more interested in thematized ads, and a better source of income for advertisers.

4.1.2. Customer care

Usually medium- and big-sized companies, with customers support, handle a large amount of unstructured information available as text documents. Relevant examples are emails, support forum discussions, documentation, shipment address information, credit card transfer reports, phone conversation transcripts, and so on. The ability of analyzing these documents and extracting the main concepts associated with them provides several concrete advantages. First of all, documents can be classified in a more effective fashion and this makes their retrieval easier. In addition, once the concepts present in a collection of documents have been extracted, it is possible to identify relevant associations between documents on the basis of the concepts they share. Ultimately, this enables to perform sophisticated data analysis targeted at discovering trends or, in case of forms, hidden associations among the products/services offered by a brand. In this context, Web Data Extraction techniques play a key role because they are required to quickly process large collections of textual documents and derive the information located in these documents. The retrieved data are finally processed by means of algorithms generally coming from the area of Natural Language Processing.

4.1.3. Database building

In the Web marketing sector, Web Data Extraction techniques can be employed to gather data referring to a given domain. These data may have a twofold effect: (i) a design, through reverse engineering analysis, can design and implement a DBMS representing that data; (ii) the DBMS can be automatically populated by using data provided by the Web Data Extraction system. Activities (i) and (ii) are also called *Database Building*. Fields of application of Database Building are countless: financial companies could be interested in extracting financial data from the Web and storing them in their DBMSs. Extraction tasks are often scheduled so that to be executed automatically and periodically. Also the real estate market is very florid: acquiring data from multiple Web sources is an important task for a real estate company, for comparison, pricing, co-offering, etc. Companies selling products or services probably want to *compare their pricing with other competitors*: products pricing data extraction is an interesting application of Web Data Extraction systems. Finally we can list other related tasks, obviously involved in the Web Data Extraction: duplicating an on-line database, extracting dating sites information, capturing auction information and prices from on-line auction sites, acquiring job postings from job sites, comparing betting information and prices, etc.

4.1.4. Software Engineering

Extracting data from Web sites became interesting also for Software Engineering: for instance, Rich Internet Applications (RIAs) are rapidly emerging as one of the most innovative and advanced kind of application on the Web. RIAs are Web applications featuring a high degree of interaction and usability, inherited from the similarity to desktop applications. Amalfitano et al. [2] have developed a reverse engineering approach to abstract Finite States Machines representing the *client-side behavior* offered by RIAs.

4.1.5. Business Intelligence and Competitive Intelligence

Baumgartner et al. [9, 12, 10] deeply analyzed how to apply Web Data Extraction techniques and tools to improve the process of acquiring market information. A solid layer of knowledge is fundamental to optimize the decision-making activities and a large amount of public information could be retrieved on the Web. They illustrate how to acquire these unstructured and semi-structured information. In particular, using the Lixto Suite to access, extract, clean and deliver data, it is possible to gather, transform and obtain information useful to business purposes. It is also possible to integrate these data with other common platforms for Business Intelligence, like SAP (www.sap.com) or Microsoft Analysis Services [88].

Wider, the process of gathering and analyzing information about products, customers, competitors with the goal of helping the managers of a company in decisional processes is commonly called *Competitive Intelligence*, and is strictly related to data mining [58]. Zanasi [125] was the first to introduce the possibility of acquiring these data, through data mining processes, on public domain information. Chen et al. [22] developed a platform, that works more like a spider than a Web Data Extraction system, which represents a useful tool to support Competitive Intelligence operations. In Business Intelligence scenarios, we ask Web Data Extraction techniques to satisfy two main requirements: scalability and efficient planning strategies because we need to extract as much data as possible with the smallest amount of resources in time and space.

4.1.6. *Web process integration and channel management*

In the Web of today data is often available via APIs (e.g., refer to: www.programmableweb.com). Nevertheless, the larger amount of data is primarily available in semi-structured formats such as HTML. To use Web data in Enterprise Applications and service-oriented architectures, it is essential to provide means for automatically turning Web Applications and Web sites into Web Services, allowing a structured and unified access to heterogeneous sources. This includes to understand the logic of the Web application, to fill out form values, and to grab relevant data.

In a number of business areas, Web applications are predominant among business partners for communication and business processes. Various types of processes are carried out on Web portals, covering activities such as purchase, sales, or quality management, by manually interacting with Web sites. Typical vertical examples, where Web Data Extraction proves useful include channel management in the travel industry (like automating the regular offerings of rooms on hotel portals with bi-directional Web connectors), re-packaging complex Web transactions to Web services and consequently to other devices, as well as automating communication of automotive suppliers with automotive companies. Tools for wrapper generation pave the way for *Web Process Integration* and enable the *Web of Services*, i.e., the seamless integration of Web applications into a corporate infrastructure or service oriented landscape by generating Web services from given Web sites [5]. Web process integration can be understood as front-end and “outside-in” integration: integrate cooperative and non-cooperative sources without the need for information provider to change their back-end. Additional requirements in such scenarios include to support a large number of users and real-time parametrized Web queries and support of complex Web transactions.

4.1.7. *Functional Web application testing*

Testing and Quality Management are essential parts of the life-cycle of software. Facets of testing are manifold, including functional tests, stress/load tests, integration tests, and testing against specifications, to name a few. Usually, the strategy is to automate a large percentage of functional tests and execute test runs as part of nightly builds as regression tests. Such tests occur at various levels, for example testing the system functionality as a blackbox via APIs, or testing the system at the GUI level simulating either the user’s steps or creating a model of possible application states.

In today’s world of Software-as-a-Service platforms and Web oriented architectures, Web application testing plays an important role. One aspect is simulating the user’s path through the application logic. Robust identification criteria can be created by taking advantage of the tree structure or visual structure of the page. Typical actions in such test scripts include to set/get values of form fields, picking dates, checkpoints to compare values, and following different branches depending on the given page. Due to automation, every step can be parametrized and a test script executed in variations. The requirements for tools in the area of Web application testing are to deal well with AJAX/dynamic HTML, to create robust test scripts, to efficiently maintain test scripts, to execute test runs and create meaningful reports, and, unlike other application areas, the support of multiple state-of-the-art browsers in various versions is an absolute must. One widely used open source tool for Web application testing is *Selenium* (seleniumhq.org).

4.1.8. Comparison shopping

One of the most appreciated services in e-commerce area is the *comparison shopping*, i.e., the capability of comparing products or services; various type of comparisons are allowed going from simple prices comparison to features comparison, technical sheets comparison, user experiences comparison, etc. These services heavily rely on Web Data Extraction, using Web sites as sources for data mining and a custom internal engine to make possible the comparison of similar items. Many Web stores today also offer *personalization forms* that make the extraction tasks more difficult: for this reason many last-generation commercial Web Data Extraction systems (e.g., Lixto, Kapow Mashup Server, UnitMiner, Bget) provide support for deep navigation and dynamic content pages.

4.1.9. Mashup scenarios

Today, leading software vendors provide mashup platforms (such as Yahoo! Pipes or Lotus Mashups) and establish mashup communication standards such as EMMML (see: www.openmashup.org). A mashup is a Web site or a Web application that combines a number of Web sites into an integrated view. Usually, the content is taken via APIs, embedding RSS or Atom Feeds in a REST-like way. With wrapper technology, one can leverage legacy Web applications to light-weight APIs such as REST that can be integrated in mashups in the same fashion. Web Mashup Solutions no longer need to rely on APIs offered by the providers of sites, but can extend the scope to the whole Web. In particular, the deep Web gets accessible by encapsulating complex form queries and application logic steps into the methods of a Web Service. End users are put in charge of creating their own views of the Web and embed data into other applications (“consumers as producers”), usually in a light-weight way. This results in “situational applications”, possibly unreliable and unsecure applications, that however help to solve an urgent problem immediately. In Mashup scenarios, one important requirement of Web Data Extraction tools is the ease of use for non-technical content managers, to give them the possibility to create new Web connectors without help of IT experts.

4.1.10. Opinion mining

Related to comparison shopping, the opinion sharing represents its evolution: users want to *express opinions* on *products, experiences, services* they enjoyed, etc. The most common form of opinion sharing is represented by blogs, containing articles, reviews, comments, tags, polls, charts, etc. As an example, we can cite MOpiS (Multiple Opinion Summarizer) [71], an algorithm that generates summaries of reviews associated with commercial products by taking into account both the review content and some metadata (e.g., the usefulness of a review, the technical expertise of the reviewer and so on) that are usually present together with review text. All this information usually lacks of structure, so their extraction is a big problem, also for current systems, because of the billions of Web sources now available. Sometimes model-based tools fit good, taking advantage of common templates (e.g., Wordpress (www.wordpress.org), Blogger (www.blogger.com), etc.), other times Natural Language Processing techniques fit better. Kushal et al. [32] approached the problem of opinion extraction and subsequent semantic classification of reviews of products.

Another form of opinion sharing in semi-structured platforms is represented by Web portals that let users to write unmoderated opinions on various topics.

4.1.11. Citation databases

Citation databases building is an intensive Web Data Extraction field of application: CiteSeer (citeseer.ist.psu.edu), Google Scholar and DBLP (www.informatik.uni-trier.de/ley/db), amongst others, are brilliant examples of applying Web Data Extraction to approach and solve the problem of collect digital publications, extract relevant data, i.e., *references and citations*, and build structured databases, where users can perform searches, comparisons, count of citations, cross-references, etc.

Several challenges are related to this context of application: for example, the corpus of scientific publications could rapidly vary over time and, to keep the database updated it could be necessary to repeatedly apply a Web Data Extraction process *from scratch* on the same Web sources. Such an operation, however, can be excessively time-consuming. An attempt to address this challenge has been done in [21]: in that paper, the authors suggest an *incremental solution* which requires to identify portions of information shared by consecutive snapshots and to reuse the information extracted from a snapshot to the subsequent one.

4.1.12. Web accessibility

Techniques for automatic data extraction and document understanding are extremely helpful in making Web pages more accessible to blind and partial-sighted users.

Today's solution approaches are inefficient to overcome the problem. The first approach, screen-reader usage, is optimized for native client interfaces, and not well equipped to deal with the presentation, content and interactions in Web 2.0 - such as understanding the reading order, telling the user what a date picker means, or jumping from one forum post to the next. The second approach, the Web Accessibility Initiative, is no doubt absolutely necessary and has defined valuable concepts, such as ARIA roles assignable to GUI elements. However, due to the additional investment at best such guidelines are applied in governmental sites.

Approaches such as ABBA [37] overcome these limitations. In the ABBA approach, a Web page is transformed into a formal multi-axial semantic model; the different axes offer means to reason on and serialize the document by topological, layout, functional, content, genre and saliency properties. A blind person can navigate along and jump between these axes to skip to the relevant parts of a page. E.g., the presentational axis contains transformed visual cues, allowing the user to list information in the order of visual saliency.

4.1.13. Main content extraction

Typical Web pages, for example news articles, contain, additionally to the main content, navigation menus, advertisements and templates. In some cases, such as when archiving a news article Web page for later offline reading it is convenient to get rid of such irrelevant fragments. To extract the main content only, one needs to apply techniques to distinguish the relevant content from the irrelevant one. Approaches range from complex visual Web page analysis to approaches leaning on text density or link density analysis. An approach for boilerplate detection using shallow text features was introduced in [70]². Additionally, tools/apps such as InstaPaper or the Readability Library – The reader is referred to the PHP port www.keyvan.net/2010/08/php-readability – use main content extraction to store the relevant fragment and text from a Web page that resembles the article for later reading.

4.1.14. Web (experience) archiving

Digital preservation and Web data curation are the goals of the discipline of Web Archiving. On the one hand, this means to access information no longer available in the live Web, and on the other hand also to reflect how the Web was used in former times. Events such as iPres and IAWW, and consortia such as Netpreserve, as well as various local Web archiving initiatives, tackle this task. There are numerous challenges [85] due to the facts that Web pages are ephemeral and due to unpredictable additions, deletions and modifications. Moreover, the hidden Web poses a further challenge. Approaches to Web Archiving are manifold, and range from Web Crawling/Harvesting, server-side archiving, transaction-based archiving, archiving the content of Web databases to library-like approaches advocating persistent identifiers (e.g., Digital Object Identifier - DOI). Especially in the use case about Web database content archiving, Web data extraction techniques are exploited.

²The associated software can be found at <http://code.google.com/p/boilerpipe/>

Another possibility of archiving the Web is to archive how Web pages are consumed. We can consider this as a type of event-based archiving, that makes sense especially for rich Web applications. The idea is not to archive everything, but to archive selectively sample pathes through an application. This requires to choose sample sites, get an understanding about common and frequent pathes through an application and store the interaction sequence. In a museum-like approach, such selected sequences are stored and can be restored or replayed to provide the user with the experience how the Web was consumed. Deep Web navigation techniques and form understanding are the key technologies here. In the Social Web scenario, the Blogforever project³ developed some techniques to harvest, preserve, manage and reuse blog content. The information extraction procedure is fully automated and uses Web feeds (i.e., structured XML documents allowing to access the content of a Web page) plus the hypertext of the blog [53]. Web feeds are used to train and generate a wrapper which is described by a set of simple extraction rules. Retrieved content is finally parsed in such a way as to render the captured content into structured XML data.

4.1.15. Summary

Figure 9 summarizes the 14 discussed enterprise application scenarios. For each scenario we describe the main ingredients of the value chain, as illustrated in the scenario descriptions.

4.2. Social Web Applications

In the latest years, Social Web platforms emerged as one of the most relevant phenomenon on the Web: these platforms are built around users, letting them to create a *web of links between people*, to share thoughts, opinions, photos, travel tips, etc. In such a scenario, often called *Web 2.0* users turn from *passive consumers* of contents to *active producers*. Social Web platforms provide novel and unprecedented research opportunities. The analysis on a large, often planetary, scale of patterns of users interactions provides the concrete opportunity of answering questions like these: how does human relationships (e.g., friendship relationships) are created and evolve over time [69]? How does novel ideas spread and propagate through the web of human contacts [15]? How does the human language evolve through social interactions (e.g., how do person expand their lexicon on the basis of their interactions with other persons) [86]?

Besides scientific questions, the analysis of patterns of human interactions in Social Web platforms has also relevant implication at the business level: if we are able to understand the dynamics of interactions among humans, we are also able to identify how groups of users aggregate themselves around shared interests. This is a crucial step for marketing purposes: once users have been grouped, we can, for instance, selectively disseminate commercial advertisements only to those groups formed by users who are actually interested in receiving those advertisements. In an analogous fashion, the fabric of social interactions can be used to identify *influential users*, i.e., those users whose commercial behaviors are able to stimulate the adoption/rejection of a given product by large masses of users.

Finally, Social Web users often create accounts and/or profiles in *multiple platforms* [112, 34]. Correlating these accounts and profiles is a key step to understand how the design features and the architecture of a Social Web platform impact on the behavior of a user: so, for instance, one may ask whether some functionalities provided by a given platform augment the aptitude of users to socialize or they impact on the volume of contents produced by a user. Once the relationship between the features of a given platform and the behavior of a user has been elucidated, the designers/managers of that platform can provide novel services to raise the level of engagement of a user in the platform or to raise their degree of loyalty (e.g., to avoid users become inactive in the platform and migrate to other ones).

In the context above, Web Data Extraction techniques play a key role because the capability of timely gathering large amounts of data from one or more Social Web platforms is an indefeasible tool to analyze

³blogforever.eu

Context-Aware Advertising	NER/NLP	Classification				
Customer Care	Search	NER/NLP	Classification			
Database Building	Search	Fact Extraction	Schema Matching			
Software Engineering	Libraries					
Competitive Intelligence	Deep Web Processes	Fact Extraction	Data Cleansing and Integration	Record Linkage	Data Mining	Reports/ Analytics
Web Process Integration	Deep Web Processes	Fact Extraction	Bi-directional Enterprise integration			
Web Application Testing	Deep Web Processes	Comparisons and Assertions	Test Management			
Comparison Shopping	Deep Web Processes	Fact Extraction	Process Repackaging			
Mashups	Fact Extraction	Light-weight APIs				
Opinion Mining	Search	NER/NLP	Data Mining			
Citation Databases	Fact Extraction	Data Integration				
Web Accessibility	Fact Extraction	Content Repackaging	Process Repackaging			
Main Content Extraction	Text Extraction					
Web Experience Archiving	Search	Classification	Deep Web Processes	Experience Archiving		

Figure 9: Summary of the application domains in the enterprise context.

human activities. Traditional Web Data Extraction techniques are challenged by new and hard problems both at the technical and scientific level. First of all, Social Web platforms feature a high level of dynamism and variability because the web of contacts of a user or a group of users may significantly vary in small time slots; therefore, we need to design Web Data extraction algorithms/procedures capable of gathering large amount of data in a quick fashion, so that the fragments of collected data keep pace with changes occurring in the structure of the user social network. If such a requirement is not satisfied, the picture emerging from the analysis we can carry out on the data at our disposal could be wrong and it would fail to capture the structure and evolution of human interactions in the platform. A second challenge depends on the fact that Web Data Extraction algorithms are able to capture only a portion of the data generated within one or more platforms. Therefore, we are required to check that the features of a data sample generated as the output of a Web Data Extraction algorithm replicate fairly well the structure of the original data within the platform(s). Finally, since data to gather are associated with humans or reflect human activities, Web Data Extraction techniques are challenged to provide solid guarantees that user privacy is not violated.

In the remaining of this section we will provide an answer to these questions and illustrate how Web Data Extraction techniques have been applied to collect data from Social Web platforms. First of all, we describe how to gather data about user relationships and activities within a single social platform (see Section 4.2.1). Secondly, we consider users who created and maintain an account in multiple Social Web platforms and discuss issues and challenges related to data collection in this scenario (see Section 4.2.2).

4.2.1. Extracting data from a single Online Social Web platform

In this section we first discuss the technological challenges arising if we aim at collecting data about *social relationships*, *user activities* and the *resources* produced and shared by users. Subsequently, we discuss privacy risks associated with the consumption/usage of human related data.

Technical challenges for collecting data from a single Social Web platform. We can classify techniques to collect data from a Social Web platform into two main categories: the former category relies on the usage of ad-hoc APIs, usually provided by the Social Web platform itself; the latter relies on HTML scraping.

As for the first category of approaches, we point out that, today, Social Web platforms provide powerful APIs (often available in multiple programming languages) allowing to retrieve in an easy and quick fashion a wide range of information from the platform itself. This information, in particular, regards not only social connections involving members of the platforms but also the content the users posted and, for instance, the tags they applied to label available content.

We can cite the approach of [78] as a relevant example of how to collect data from a Social Web platform by means of an API. In that paper, the authors present the results of the crawling of the whole Twitter platform. The dataset described in [78] built consisted of 41.7 million user profiles, 1.47 billion social relations; in addition to collecting information about user relationships, the authors gathered also information on tweets and, by performing a semantic analysis, also on the main topics discussed in these tweets. The final dataset contained 4,262 trending topics and 106 million tweets.

From a technical standpoint we want to observe that: (i) The Twitter API allows to access the whole *social graph*, i.e., the graph representing users and their connections in Twitter, without authentication. Other Social Web platforms and the APIs they offer, however, do not generally allow to access the whole social graph. A meaningful example is given by the Facebook API. (ii) The Twitter API, by default, allows a human user or a software agent to send only 150 requests per hour: this could be an inadmissible limitation because the amount of information generated within Twitter in a relatively small time slot can be very large, and, then, changes in the Twitter network topology could not be properly sensed. To overcome this problem, Twitter offers *white lists*: users registered to white lists can send up to 20 000 requests per IP per hour. In [78] the authors used a group of 20 computers, each of them belonging to the Twitter white lists, to perform a real-time monitoring of Twitter.

Approaches based on the scraping of HTML pages are able to overcome the limitations above even if they are more complicated to design and implement. To the best of our knowledge, one of the first attempt to crawl large Online Social Networks was performed by Mislove et al. [90]. In that paper, the authors focused on platforms like Orkut, Flickr and LiveJournal. To perform crawling, the approach of [90] suggests to iteratively retrieve the list of friends of a user which have not yet been visited and to add these contact to the list of users to visit. According to the language of the graph theory, this corresponds to perform a Breadth-First-Search (BFS) visit of the Social Network graph. The user account from which the BFS starts is often called *seed node*; the BFS ends when the whole graph is visited or, alternatively, a stop criterium is met. The BFS is easy to implement and efficient; it produces accurate results if applied on social graphs which can be modeled as unweighted graphs. Due to these reasons, it has been applied in a large number of studies about the topology and structure of Online Social Networks (see, for instance, [20, 119, 52, 123, 17]).

As observed by [90], BFS may incur in heavy limitations. First of all, a crawler can get trapped into a strong connected component of the social graph. In addition, if we would use the BFS sample to estimate some structural properties of the social network graph, some properties could be overestimated while others could be underestimated [74].

To alleviate these problems, several authors suggested more refined sampling techniques. The implementation of these techniques is equivalent to define new Web Data Extraction procedures. Most of them have been discussed and exploited in the context of Facebook [52, 17] but, unfortunately, some of them can not be extended to other platforms.

In particular, Gjoka et al. [52] considered different visiting algorithms, like BFS, “Random Walks” and “Metropolis-Hastings Random Walks”. A particular mention goes to a visiting method called *rejection sampling*. This technique relies on the fact that a truly uniform sample, of Facebook users can be obtained by generating uniformly at random a 32-bit user ID and, subsequently, by polling Facebook about the existence of that ID. The correctness of this procedure derives from the fact that each Facebook user was uniquely identified by a numerical ID ranging from 0 and $2^{32} - 1$. Of course, such a solution works well for Facebook but it could not work for other platforms.

In [17], the authors designed a Web Data Extraction architecture based on Intelligent Agents (see Figure 1 of [17]). Such an architecture consists of three main components: (i) a server running the mining agent(s); (ii) a cross-platform Java application, which implements the logic of the agent; (iii) an Apache interface, which manages the information transfer through the Web. The proposed architecture is able to implement several crawling strategies like BFS or rejection sampling.

The sampling procedure in [17] works as follows: an agent is activated and it queries the Facebook server(s) to obtain the list of Web pages representing the list of friends of a Facebook user. Of course, the Facebook account to visit depends on the basis of the crawling algorithm we want to implement. After parsing this list of pages, it is possible to reconstruct a portion of the Facebook network. Collected data can be converted in XML format so that they can be exploited by other applications (e.g., network visualization tools).

In addition to gathering data about social relationship, we may collect contents generated by users. These contents may vary from resources posted by users (like photos in Flickr or videos in YouTube) to tags applied for labeling resources with the goal of making the retrieval of these resources easier or to increase the visibility of the contents they generated.

Privacy pitfalls in collecting user related data. The major concern about the extraction of data from Social Web platforms is about *privacy*. Several researchers, in fact, showed that we can disclose private information about users by leveraging on publicly available information available in a Social Web platform. For instance, [59] provided an approach to finding user location on the basis of user tweets. The authors used basic machine learning techniques which combined user tweets with geotagged articles in Wikipedia. In an analogous fashion, [67] used geographic coordinates extracted from geotagged Twitter data to model user location. Locations were modeled at different levels of granularity ranging from the zip code to the country level. Experimental studies in [67] show that the proposed model can predict country, state and city with an accuracy comparable with that achieved by industrial tools for geo-localization.

Crandall et al. [26] investigated how to organize a large collection of geotagged photos extracted from Flickr. The proposed approach combined content analysis (based on the textual tags typed by users to describe photos), image analysis and structural analysis based on geotagged data. The most relevant result is that it is possible to locate Flickr photos with a high precision by identifying landmarks via visual, temporal and textual features.

In [18] the authors considers a range of user signals expressing user interests (e.g., “I like” declaration in Facebook) which are often disclosed by users. By means of a semantic-driven inference technique based on an ontologized version of Wikipedia, the authors show how to discover hidden information about users like gender, relationship status and age.

From the discussion above it emerges that collecting seemingly innocuous data from a Social Web platform hides high privacy risks. Therefore, collected data should be manipulated so that reducing these risks.

Finally, due to privacy settings, the so-called *black hole problem* can arise [123]. In detail, some OSN users may decide to hide their personal information (like contact lists and posts) to strangers and, in this way, they obstacle the crawler in collecting OSN data. [123] experimentally studies the impact of black holes by considering both different crawling strategies as well as different OSNs (like YouTube and LiveJournal). Black holes have a small *but non-negligible* impact on the number of vertices and edges visited by the crawler:

on average, vertex/edge coverage decreases in fact between 7% and 9% w.r.t. a configuration in which black holes are not present.

4.2.2. Extracting data from multiple Online Social Web platforms

Social Web users often create and maintain different profiles in different platforms with different goals (e.g., to post and share their photos in Flickr, to share bookmarks on Delicious, to be aware on job proposals on LinkedIn and so on). We discuss technical approaches to collecting data from multiple Social Web platforms as well as the opportunities coming from the availability of these data. Finally, we describe potential privacy risks related to the management of gathered data.

Collecting data from multiple Social Web platforms. The main technical challenges encountered by Web Data Extraction techniques to collect data from multiple Social Web platforms consists of linking information referring to the same user or the same object. Early approaches were based on “ad-hoc” techniques; subsequent approaches featured a higher level of automation and were based on ad-hoc APIs like the Google Social Graph API.

To the best of our knowledge, one of the first approaches facing the problem of correlating multiple user accounts was presented in [112]. In that paper, the authors started with a list of 667, 141 user accounts on Delicious such that each account was uniquely associated with a Delicious profile. The same procedure was repeated on a data sample extracted from Flickr. The first stage in the correlation process consisted in comparing usernames in Delicious and Flickr: if these strings exactly match, the two accounts were considered as referring to the same person. In this way, it was possible to build a *candidate list* consisting of 232, 391 usernames such that each user name referred to a Flickr and Delicious profile. Of course, such a list must be refined because, for instance, different users may choose the same username. Since both in Delicious and Flickr the users had the chance of filling a form by specifying their real names, the authors of [112] suggested to refine the candidate list by keeping only those user accounts whose real names matched exactly. Such a procedure significantly lowered the risk of incurring in false associations but at the same time, dramatically reduced the size of the candidate list to only 502 elements. Some tricks aiming at producing a larger and accurate dataset were also proposed in [112]: in particular, the authors observed that, in real scenarios, if a user creates accounts in several Web sites, she frequently adds a link to her accounts: so by using traditional search engines, we can find all pages linking to the homepage of a user and, by filtering these hits we can find the exact URL of the profile of a user in different platforms.

The process described above can be automatized by exploiting *ad hoc tool*. Among these tools, the most popular was perhaps the **Google Social Graph API**, even if such an API is no longer available. Such an API is able to find connections among persons on the Web. It can be queried through an HTTP request having a URL called *node* as its parameter. The node specifies the URL of a Web page of a user *u*. The Google Social Graph API is able to return two kinds of results:

- *A list of public URLs that are associated with u*; for instance, it reveals the URLs of the blog of *u* and of her Twitter page.
- *A list of publicly declared connections among users*. For instance, it returns the list of persons who, in at least one social network, have a link to a page which can be associated with *u*.

The Google Social Graph API has been used in [1].

The identification of the connections between persons or Web objects (like photos or videos) is the key step to design advance services often capable of offering a high level of personalization. For instance, in [108], the authors suggest to use tags present in different Social Web systems to establish links between items located in each system. In [1] the system Mypes is presented. Mypes supports the linkage, aggregation, alignment and semantic enrichment of user profiles available in various Social Web systems, such as Flickr, Delicious and Facebook. In the field of Recommender Systems, the approach of [34] show how to merge

ratings provided by users in different Social Web platforms with the goal of computing reputation values which are subsequently used to generate recommendations.

Some authors have also proposed to combine data coming from different platforms but referring to the same object. A nice example has been provided in [111]; in that paper the authors consider users of blogs concerning music and users of Last.fm, a popular folksonomy whose resources are musical tracks. The ultimate goal of [111] is to enrich each Social Web system by re-using tags already exploited in other environments. This activity has a twofold effect: it first allows the automatic annotation of resources which were not originally labeled and, then, enriches user profiles in such a way that user similarities can be computed in a more precise way.

Privacy risks related to the management of user data spread in multiple platforms. The discussion above show that the combination and linkage of data spread in independent Social Web platforms provide clear advantages in term of item description as well as on the level of personalization that a system is able to provide to its subscribers. However, the other side of the coin is given by the *privacy problems* we may incur when we try to glue together data residing on different platforms.

Some authors recently introduced the so-called *user identification* problem, i.e., they studied what information is useful to disclose links between the multiple accounts of a user in independent Social Web platforms.

One of the first approaches to dealing with the user identification problem was proposed in [115]. In that paper the authors focused on Facebook and StudiVZ (studivz.net) and investigated which profile attributes can be used to identify users. In [124], the authors studied 12 different Social Web systems (like Delicious, Flickr and YouTube) with the goal of finding a mapping involving the different user accounts. This mapping can be found by applying a traditional search engine. In [62], the authors suggest to combine profile attributes (like usernames) with an analysis of the user contributed tags to identify users. They suggest various strategies to compare the tag-based profiles of two users and some of these strategies were able to achieve an accuracy of almost 80% in user identification. [97] explored the possibility of linking users profiles only by looking at their usernames. Such an approach is based on the idea that the probability that two usernames are associated with the same person depends on the entropies of the strings representing usernames.

[4] described a simple but effective attack to discover the multiple digital identities of a user. Such an attack depends on the fact that a user often subscribes to multiple Social Web platforms by means of a *single e-mail address*. An attacker can query Social Web platforms by issuing a list of e-mail addresses to each Social Web platform; once the profiles of a user have been identified in each platform the data contained in these profiles can be merged to obtain private personal information. The authors considered a list of about 10.4 millions of e-mail addresses and they were able to automatically identify more than 1.2 millions of user profiles registered in different platforms like Facebook and XING. The most popular providers acknowledged the privacy problem raised in [4] and implemented ad hoc countermeasures.

An interesting study is finally presented in [54]. The authors showed that correlating seemingly innocuous attributes describing user activities in multiple social networks allow attackers to track user behaviors and infer their features. In particular, the analysis of [54] focused on three features of online activity like the geo-location of users posts, the timestamp of posts, and the user's writing style (captured by proper language models).

4.3. Opportunities for cross-fertilization

In this section we discuss on the possibility of re-using Web Data Extraction techniques originally developed in a given application domain to another domain. This discussion is instrumental in highlighting techniques which can be applied across different application domains and techniques which require some additional information (which can be present in some application domains and missing in others).

In the most general case, no assumption is made on the structure and content of a collection of Web pages which constitute the input of a Web Data Extraction tool. Each Web page can be regarded as a text document and, in such a configuration, approaches relying on Regular Expressions can be applied independently of the application domain we are considering. As mentioned before, regular expressions can be regarded as a formal language allowing to find strings or patterns from text according to some matching criteria. Approaches based on regular expressions therefore may be suitable when we do not have, at our disposal, any information about the structure of the Web pages. Of course, these approaches can be extended so that to take some structural elements into account (like HTML tags). The usage of regular expression may be disadvantageous if need to collect large collections of documents and assume that those documents deal with different topics. In such a case, we need complex expressions for extracting data from documents and this would require a high level of expertise.

The next step consists of assuming that some information on the structure of the document is available. In such a case, wrappers are an effective solution and they are able to work fairly well in different application domains. In fact, the hierarchical structure induced by HTML tags associated with a page often provides useful information to the Web Data Extraction task. A powerful solution taking advantage of the structure of HTML pages derives from the the usage of XPath to quickly locate an element (or multiple instances of the same element in the document tree). Approaches based on XPath have been first exploited in the context of Enterprise applications and, later, they have been successfully re-used in the context of Social Web applications. There are, however, a number of requirements an application domain should satisfy in order to enable the usage of XPath. First of all, the structure of Web pages should be perfectly defined to the wrapper designer (or to the procedure inducing the wrapper); such an assumption, of course, can not be true in some domains because data are regarded as *proprietary* and technical details about the structure of the Web wrappers in that domain are not available. In addition, we should assume a certain *level of structural coherence* among all the Web pages belonging to a Web source. Such an assumption is often true both in enterprise domain and in many Social Web systems: in fact, Web pages within an organization (e.g., a company) or a Social Web platform often derive from the same template and, therefore, a certain form of structural regularity across all the pages emerge. By contrast, if we plan to manage pages with different structures (even if referring to the same domain), we observe that even small changes in the structure of two Web pages may have a devastating impact and thus would require to entirely rewrite the wrapper.

A nice example of cross-fertilization is in the context of the crawling of Online Social Networks [17]. In that paper the authors implemented a Web wrapper to crawl Facebook largely exploiting techniques and algorithms which were part of the Lixto suite and that were originally designed to work for Business and Competitive Intelligence applications. The proposed wrapper is based on the execution of XPath queries; human experts, in the configuration phase, are allowed to specify what elements have to be selected.

The crawler provides two running modes: (i) visual extraction and, (ii) HTTP request-based extraction. In the visual extraction mode, the crawler embeds a Firefox browser interfaced through XULRunner (developer.mozilla.org/en/XULRunner) via XPCOM. The visual approach requires the rendering of the Web pages which is a time-consuming activity. Therefore, to extract large amounts of data, [17] suggested to send HTTP requests to fetch Web pages.

In some application scenarios, in addition to assuming that the syntactic structure of a Web page is known, it is possible to assume that a rich semantic structure emerges from the Web pages. If such an hypothesis holds true, techniques from Information Extraction (IE) and Natural Language Processing (NLP) can be conveniently used [121, 13, 84]. The range of applications benefiting from NLP techniques comprises relevant examples both in the enterprise and Social Web scenarios: for instance, relevant applications of NLP/IE techniques are the extraction of facts from speech transcriptions in forums, email messages, newspaper articles, resumes etc.

5. Conclusions

The World Wide Web contains a large amount of unstructured data. The need for structured information urged researchers to develop and implement various strategies to accomplish the task of automatically extracting data from Web sources. Such a process is known with the name of Web Data Extraction and it has had (and continues to have) a wide range of applications in several fields, ranging from commercial to Social Web applications.

The central thread of this survey is to classify existing approaches to Web Data Extraction in terms of the applications for which they have been employed.

In the first part of this paper, we provided a classification of algorithmic techniques exploited to extract data from Web pages. We organized the material by presenting first basic techniques and, subsequently, the main variants to these techniques. Finally, we focus on how Web Data Extraction systems work in practices. We provide different perspectives to classify Web Data Extraction systems (like the ease of use, the possibility of extracting data from the Deep Web and so on).

The second part of the survey is about the applications of Web Data Extraction systems to real-world scenarios. We provided a simple classification framework in which existing applications have been grouped into two main classes (Enterprise and Social Web applications). The material has been organized around the application domains of Web Data Extraction systems: we identify, for each class, some sub-domains and described how Web Data Extraction techniques work in each domain. This part ends with a discussion about the opportunities of cross-fertilization.

We discuss now some possible future (and, in our opinion, promising) applications of Web Data Extraction techniques.

Bio-informatics and Scientific Computing. A growing field of application of Web Data Extraction is *bio-informatics*: on the World Wide Web it is very common to find medical sources, in particular regarding *bio-chemistry* and *genetics*. Bio-Informatics is an excellent example of the application of scientific computing – refer e.g. to [35] for a selected scientific computing project.

Plake et al. [99] worked on PubMed (www.pubmed.com) – the biggest repository of medical-scientific works that covers a broad range of topics – extracting information and relationships to create a graph; this structure could be a good starting point to proceed in extracting data about *proteins* and *genes*, for example connections and interactions among them: this information can be usually found, not in Web pages, rather they are available in PDF or Postscript format. In the future, Web Data Extraction should be extensively used also to these documents: approaches to solving this problem are going to be developed, inherited, both from Information Extraction and Web Data Extraction systems, because of the semi-structured format of PostScript-based files. On the other hand, web services play a dominant role in this area as well, and another important challenge is the intelligent and efficient querying of Web services as investigated by the ongoing SeCo project (www.search-computing.it).

Web harvesting. One of the most attractive future applications of the Web Data Extraction is *Web Harvesting* [118]: Gatterbauer [50] defines it as “the process of gathering and integrating data from various heterogeneous Web sources”. The most important aspect (although partially different from specific Web Data Extraction) is that, during the last phase of data transformation, the amount of gathered data is *many times greater* than the extracted data. The work of filtering and refining information from Web sources ensures that extracted data lie in the domain of interest and are relevant to users: this step is called *integration*. The Web harvesting remains an open problem with large margin of improvement: because of the billions of Web pages, it is a computational problem, also for restricted domains, to crawl enough sources from the Web to build a *solid ontological base*. There is also a human engagement problem, correlated to the degree of automation of the process: when and where humans have to interact with the

system of Web harvesting? Should be a fully automatic process? What degree of precision can we accept for the harvesting? All these questions are still open for future works. Projects such as the DIADEM (web.comlab.ox.ac.uk/projects/DIADEM) at Oxford University tackle the challenge for fully automatic generation of wrappers for restricted domains such as real estate.

Linking Data from several Web sources. A growing number of authors suggest to integrate data coming from as many sources as possible to obtain a detailed description of an object. Such a description would be hard to obtain if we would focus only on a single system/service.

The material discussed in this section partially overlaps with some ideas/techniques presented in Section 4.2.2. However, in Section 4.2.2 we focused on Social Web platforms and showed that linking information stored in the profiles of a user spread across multiple platforms leads to a better identification of user needs and, ultimately, raises the level of personalization a system can offer to her. In this section, by contrast, we focus on systems having two main features. (i) they often publicly expose their data on the Web and (ii) these systems are *lacking* of a social connotation, i.e., they do not target at building a community of interacting members.

A first research work showing the benefits of linking data provided by independent Web systems is provided in [113]. In that paper, the authors combined information from the Internet Movie Database (www.imdb.com) and Netflix (www.netflix.com). The IMDB is an online database containing extensive information on movies, actors and television shows. IMDB users are allowed to add tags to describe the main features of a movie (e.g., the most important scenes, location, genres and so on). Netflix is an US-based company offering an online DVD rental service. Netflix users are allowed to rate a movie by providing a score. In [113], data from Netflix and IMDB were imported in a relational DBMS; movie titles in IMDB were correlated with movie titles in Netflix by applying string matching techniques. In this way, for each movie, the Netflix ratings and the IMDB description were available. The authors studied three recommendation strategies based on ratings alone, on tags alone and, finally, a combination of ratings and tags. Experimental trials provided evidence that the combination of data located in different systems was able to improve the level of accuracy in the provided recommendations. Spiegel et al. presented an analogous study in [110]; in that paper the authors combine user ratings (coming from MovieLens database) with movie information provided by IMDB.

The linkage of datasets coming from independent Web platforms fuels novel scientific applications. One of these application is given from *cross-domain recommender systems*, i.e., on recommender systems running on multiple domains. For example, one can use user ratings about movies to predict user preferences in the music domain. In a given domain (e.g., if our main aim is to recommend movies), information about user preferences may be insufficient and, in many case, most of this information is missing for a non-negligible part of user population. By contrast, a relevant amount of information describing user preferences could be available (at a limited cost) in other domains. We could *transfer* this information from a domain to another one with the goal of dealing with data sparsity and improving recommendation accuracy.

Various strategies have been proposed to perform such a transfer of information and some of these strategies take advantages from models and algorithms developed in the field of *transfer of learning* [96]. For instance, some approaches use co-clustering [80], other clustering techniques in conjunction with probabilistic models [81] and, finally, other approaches project the user and item space in the various domains in a shared latent space by means of regularization techniques [129].

Scalability Issues and the usage of Cloud Computing infrastructures. . Because of the increasing complexity of Web Data Extraction procedures, some authors suggested to use Cloud computing services like Amazon EC2 to provide a high level of reliability and scalability. In cloud architectures, computing resources can be accessed as services which are paid by end users depending on the usage time or on the amount of data to process. One of the first Web Data Extraction platforms relying on a cloud architecture is Lixto [12]: in it cloud clients receive a wrapper as well as extraction parameters and they pass back retrieved data;

these data are not stored on cloud instances. Cloud architectures provide intelligent procedures to fairly (and intelligently) distribute the computational load amongst the available processing units and this yields a more efficient resource management. Due to this reasons we believe that in the next years a growing number of Web Data Extraction platforms will rely on cloud services.

Extracting Web Content into Semantic Web format. . The popularity of the open linked data initiative prompted some authors to develop systems which support the extraction of Web content and their storage in some Semantic Web format. Among these systems we cite the above mentioned DEiXTo system, which DEiXTo which to build extraction rules capable of converting extracted content in any structured format (and, among them, RDF). Other related projects are Virtuoso Sponger⁴, which generates Linked Data starting from different data sources and supports a wide range of data representation formats, Semantic Fire⁵, which allows to extract and store the data from a Web site as RDF.

References

- [1] F. Abel, E. Herder, G.-J. Houben, N. Henze, and D. Krause. Cross-system User Modeling and Personalization on the Social Web. *User Modeling and User-Adapted Interaction*, 23(2-3):169–209, 2013.
- [2] D. Amalfitano, A. Fasolino, and P. Tramontana. Reverse engineering finite state machines from rich internet applications. In *Proc. 2008 15th Working Conference on Reverse Engineering*, pages 69–73, Washington, DC, USA, 2008. IEEE Computer Society.
- [3] L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna. Four degrees of separation. *Arxiv preprint arXiv:1111.4570*, 2011.
- [4] M. Balduzzi, C. Platzer, T. Holz, E. Kirda, D. Balzarotti, and C. Kruegel. Abusing social networks for automated user profiling. In *Proc. International Symposium on Recent Advances in Intrusion Detection*, Lecture Notes in Computer Science, pages 422–441, Ottawa, Ontario, Canada, 2010. Springer.
- [5] R. Baumgartner, A. Campi, G. Gottlob, and M. Herzog. Web data extraction for service creation. *Search Computing: Challenges and Directions*, 2010.
- [6] R. Baumgartner, M. Ceresna, and G. Ledermuller. Deepweb navigation in web data extraction. In *Proc. International Conference on Computational Intelligence for Modelling, Control and Automation*, pages 698–703, Washington, DC, USA, 2005. IEEE Computer Society.
- [7] R. Baumgartner, S. Flesca, and G. Gottlob. The elog web extraction language. In *Proc. Artificial Intelligence on Logic for Programming*, pages 548–560, London, UK, 2001. Springer-Verlag.
- [8] R. Baumgartner, S. Flesca, and G. Gottlob. Visual web information extraction with lixto. In *Proc. 27th International Conference on Very Large Data Bases*, pages 119–128, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [9] R. Baumgartner, O. Frölich, G. Gottlob, P. Harz, M. Herzog, P. Lehmann, and T. Wien. Web data extraction for business intelligence: the lixto approach. In *Proc. 12th Conference on Datenbanksysteme in Büro, Technik und Wissenschaft*, pages 48–65, 2005.
- [10] R. Baumgartner, K. Fröschl, M. Hronsky, M. Pöttler, and N. Walchhofer. Semantic online tourism market monitoring. *Proc. 17th ENTER eTourism International Conference*, 2010.
- [11] R. Baumgartner, W. Gatterbauer, and G. Gottlob. Web data extraction system. *Encyclopedia of Database Systems*, pages 3465–3471, 2009.
- [12] R. Baumgartner, G. Gottlob, and M. Herzog. Scalable web data extraction for online market intelligence. *Proc. 35th International Conference on Very Large Databases*, 2(2):1512–1523, 2009.
- [13] A. Berger, V. D. Pietra, and S. D. Pietra. A maximum entropy approach to natural language processing. *Comput. Linguist.*, 22(1):39–71, 1996.
- [14] M. Berthold and D. Hand. *Intelligent Data Analysis: An Introduction*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1999.
- [15] L. Bettencourt, A. Cintrón-Arias, D. Kaiser, and C. Castillo-Chávez. The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. *Physica A: Statistical Mechanics and its Applications*, 364:513–536, 2006.
- [16] M. Califf and R. Mooney. Bottom-up relational learning of pattern matching rules for information extraction. *The Journal of Machine Learning Research*, 4:177–210, 2003.
- [17] S. Catanese, P. De Meo, E. Ferrara, G. Fiumara, and A. Provetti. Crawling facebook for social network analysis purposes. In *Proc. International Conference on Web Intelligence, Mining and Semantics*, page 52, Sogndal, Norway, 2011. ACM.

⁴<http://virtuoso.openlinksw.com/dataspace/doc/dav/wiki/Main/VirtSponger>

⁵<https://code.google.com/p/semantic-fire/>

- [18] A. Chaabane, G. Acs, and M. Kaafar. You Are What You Like! Information Leakage Through Users' Interests. In *Proc. Annual Network and Distributed System Security Symposium*, 2012.
- [19] C. Chang, M. Kayed, M. Girgis, and K. Shaalan. A Survey of Web Information Extraction Systems. *IEEE Transactions on Knowledge and Data Engineering*, 18(10):1411–1428, 2006.
- [20] D. Chau, S. Pandit, S. Wang, and C. Faloutsos. Parallel crawling for online social networks. In *Proc. 16th International Conference on the World Wide Web*, pages 1283–1284, 2007.
- [21] F. Chen, A. Doan, J. Yang, and R. Ramakrishnan. Efficient information extraction over evolving text data. In *Proc. IEEE 24th International Conference on Data Engineering*, pages 943–952, Cancun, Mexico, 2008. IEEE.
- [22] H. Chen, M. Chau, and D. Zeng. Ci spider: a tool for competitive intelligence on the web. *Decis. Support Syst.*, 34(1):1–17, 2002.
- [23] W. Chen. New algorithm for ordered tree-to-tree correction problem. *Journal of Algorithms*, 40(2):135–158, 2001.
- [24] M. Collins. A new statistical parser based on bigram lexical dependencies. In *Proc. 34th annual meeting on Association for Computational Linguistics*, pages 184–191. Association for Computational Linguistics, 1996.
- [25] M. D. Conover, C. Davis, E. Ferrara, K. McKelvey, F. Menczer, and A. Flammini. The geospatial characteristics of a social movement communication network. *PloS one*, 8(3):e55957, 2013.
- [26] D. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg. Mapping the world's photos. In *Proc. 18th International Conference on World Wide Web*, pages 761–770, Madrid, Spain, 2009. ACM.
- [27] V. Crescenzi and G. Mecca. Automatic information extraction from large websites. *J. ACM*, 51(5):731–779, 2004.
- [28] V. Crescenzi, G. Mecca, and P. Merialdo. Roadrunner: Towards automatic data extraction from large web sites. In *Proc. 27th International Conference on Very Large Data Bases*, pages 109–118, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [29] V. Crescenzi, G. Mecca, and P. Merialdo. Improving the expressiveness of roadrunner. In *SEBD*, pages 62–69, 2004.
- [30] N. Dalvi, P. Bohannon, and F. Sha. Robust Web extraction: an approach based on a probabilistic tree-edit model. In *Proc. of the ACM International Conference on Management of Data (SIGMOD 2009)*, pages 335–348, Providence, Rhode Island, USA, 2009. ACM.
- [31] N. Dalvi, R. Kumar, and M. Soliman. Automatic wrappers for large scale Web extraction. *Proc. of the VLDB Endowment*, 4(4):219–230, 2011.
- [32] K. Dave, S. Lawrence, and D. Pennock. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In *Proc. 12th international conference on World Wide Web*, pages 519–528, New York, NY, USA, 2003. ACM.
- [33] P. De Meo, E. Ferrara, F. Abel, L. Aroyo, and G.-J. Houben. Analyzing user behavior across social sharing environments. *ACM Transactions on Intelligent Systems and Technology*, 5(1):14–42, 2013.
- [34] P. De Meo, A. Nocera, G. Quattrone, D. Rosaci, and D. Ursino. Finding reliable users and social networks in a social internetworking system. In *Proc. International Database Engineering and Applications Symposium*, pages 173–181, Cetraro, Cosenza, Italy, 2009. ACM Press.
- [35] M. Descher, T. Feilhauer, T. Ludescher, P. Masser, B. Wenzel, P. Brezany, I. Elsayed, A. Wöhrer, A. M. Tjoa, and D. Huemer. Position paper: Secure infrastructure for scientific data life cycle management. In *ARES*, pages 606–611, 2009.
- [36] A. Doan, J. Naughton, R. Ramakrishnan, A. Baid, X. Chai, F. Chen, T. Chen, E. Chu, P. DeRose, B. Gao, et al. Information extraction challenges in managing unstructured data. *ACM SIGMOD Record*, 37(4):14–20, 2009.
- [37] R. Fayzrahmanov, M. Goebel, W. Holzinger, B. Kruepl, A. Mager, and R. Baumgartner. Modelling web navigation with the user in mind. In *Proc. 7th International Cross-Disciplinary Conference on Web Accessibility*, 2010.
- [38] E. Ferrara. A large-scale community structure analysis in facebook. *EPJ Data Science*, 1(9):1–30, 2012.
- [39] E. Ferrara and R. Baumgartner. Automatic wrapper adaptation by tree edit distance matching. *Combinations of Intelligent Methods and Applications*, pages 41–54, 2011.
- [40] E. Ferrara and R. Baumgartner. Design of automatically adaptable web wrappers. In *Proc. 3rd International Conference on Agents and Artificial Intelligence*, pages 211–217, 28-30 January 2011.
- [41] E. Ferrara and R. Baumgartner. Intelligent self-repairable web wrappers. *Lecture Notes in Computer Science*, 6934:274–285, 2011.
- [42] E. Ferrara, M. JafariAsbagh, O. Varol, V. Qazvinian, F. Menczer, and A. Flammini. Clustering memes in social media. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 548–555. IEEE/ACM, 2013.
- [43] E. Ferrara, O. Varol, F. Menczer, and A. Flammini. Traveling trends: social butterflies or frequent fliers? In *Proceedings of the first ACM conference on Online social networks*, pages 213–222. ACM, 2013.
- [44] G. Fiumara. Automated information extraction from web sources: a survey. *Proc. of Between Ontologies and Folksonomies Workshop*, pages 1–9, 2007.
- [45] S. Flesca, G. Manco, E. Masciari, E. Rende, and A. Tagarelli. Web wrapper induction: a brief survey. *AI Communications*, 17(2):57–61, 2004.
- [46] D. Freitag. Machine learning for information extraction in informal domains. *Machine learning*, 39(2):169–202, 2000.
- [47] T. Furche, G. Gottlob, G. Grasso, O. Gunes, X. Guo, A. Kravchenko, G. Orsi, C. Schallhart, A. J. Sellers, and C. Wang. DIADEM: domain-centric, intelligent, automated data extraction methodology. In *WWW (Companion Volume)*, pages 267–270, 2012.
- [48] T. Furche, G. Gottlob, G. Grasso, C. Schallhart, and A. Sellers. OXPath: A Language for Scalable, Memory-efficient Data Extraction from Web Applications. *Proceedings of the VLDB Endowment*, 4(11):1016–1027, 2011.
- [49] J. Garrett. AJAX: A New Approach to Web Applications. Technical report, Adaptive Path, 2005.

- [50] W. Gatterbauer. Web harvesting. *Encyclopedia of Database Systems*, pages 3472–3473, 2009.
- [51] W. Gatterbauer and P. Bohunsky. Table extraction using spatial reasoning on the css2 visual box model. In *Proc. 21st national conference on Artificial intelligence*, pages 1313–1318. AAAI Press, 2006.
- [52] M. Gjoka, M. Kurant, C. Butts, and A. Markopoulou. Walking in Facebook: a case study of unbiased sampling of OSNs. In *Proc. 29th Conference on Information Communications*, pages 2498–2506. IEEE, 2010.
- [53] G. Gkotsis, K. Stepanyan, A. Cristea, and M. Joy. Self-supervised automated wrapper generation for weblog data extraction. In *Proc. of the British National Conference on Databases (BNCOD 2013) - Big Data*, pages 292–302, Oxford, United Kingdom, 2013. Springer.
- [54] O. Goga, H. Lei, S. Parthasarathi, G. Friedland, R. Sommer, and R. Teixeira. On exploiting innocuous user activity for correlating accounts across social network sites. Technical report, ICSI Technical Reports - University of Berkeley, 2012.
- [55] G. Gottlob and C. Koch. Logic-based web information extraction. *SIGMOD Rec.*, 33(2):87–94, 2004.
- [56] G. Gottlob and C. Koch. Monadic datalog and the expressive power of languages for web information extraction. *J. ACM*, 51(1):74–113, 2004.
- [57] B. Hammersley. *Developing feeds with rss and atom*. O’Reilly, 2005.
- [58] J. Han and M. Kamber. *Data mining: concepts and techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.
- [59] B. Hecht, L. Hong, B. Suh, and E. Chi. Tweets from Justin Bieber’s heart: the dynamics of the location field in user profiles. In *Proc. International Conference on Human Factors in Computing Systems*, pages 237–246, Vancouver, British Columbia, Canada, 2011. ACM.
- [60] C.-N. Hsu and M.-T. Dung. Generating finite-state transducers for semi-structured data extraction from the web. *Inf. Syst.*, 23(9):521–538, 1998.
- [61] G. Huck, P. Fankhauser, K. Aberer, and E. Neuhold. JEDI: Extracting and synthesizing information from the web. In *Proc. 3rd IFICIS International Conference on Cooperative Information Systems*, 1998.
- [62] T. Iofciu, P. Fankhauser, F. Abel, and K. Bischoff. Identifying Users Across Social Tagging Systems. In *Proc. 5th International Conference on Weblogs and Social Media*, Barcelona, Spain, 2011.
- [63] U. Irmak and T. Suel. Interactive wrapper generation with minimal user effort. In *Proc. 15th International Conference on World Wide Web*, pages 553–563, Edinburgh, Scotland, 2006. ACM.
- [64] K. Kaiser and S. Miksch. Information extraction. a survey. Technical report, E188 - Institut für Softwaretechnik und Interaktive Systeme; Technische Universität Wien, 2005.
- [65] R. Khare and T. Çelik. Microformats: a pragmatic path to the semantic web. In *Proc. 15th international conference on World Wide Web*, pages 865–866, New York, NY, USA, 2006. ACM.
- [66] Y. Kim, J. Park, T. Kim, and J. Choi. Web information extraction by html tree edit distance matching. In *International Conference on Convergence Information Technology*, pages 2455–2460. IEEE, 2007.
- [67] S. Kinsella, V. Murdock, and N. O’Hare. I’m eating a sandwich in Glasgow: modeling locations with tweets. In *Proc. International Workshop on Search and Mining user-generated contents*, pages 61–68, Glasgow, UK, 2011. ACM.
- [68] J. Kleinberg. The small-world phenomenon: an algorithm perspective. In *Proc. ACM symposium on Theory of Computing*, pages 163–170, Portland, Oregon, USA, 2000. ACM.
- [69] J. Kleinberg. The convergence of social and technological networks. *Communications of the ACM*, 51(11):66–72, 2008.
- [70] C. Kohlschütter, P. Fankhauser, and W. Nejdl. Boilerplate detection using shallow text features. In *Proc. 3rd ACM international conference on Web search and data mining*, pages 441–450. ACM, 2010.
- [71] F. Kokkoras, E. Lampridou, K. Ntonas, and I. Vlahavas. MOpiS: A multiple opinion summarizer. In *Proc. of the 5th Hellenic Conference on Artificial Intelligence (SETN 2008)*, pages 110–122, Syros, Greece, 2008. Springer.
- [72] B. Krüpl, M. Herzog, and W. Gatterbauer. Using visual cues for extraction of tabular data from arbitrary html documents. In *Proc. 14th international conference on World Wide Web*, pages 1000–1001, New York, NY, USA, 2005. ACM.
- [73] B. Krüpl-Sypien, R. R. Fayzrakhmanov, W. Holzinger, M. Panzenböck, and R. Baumgartner. A versatile model for web page representation, information extraction and content re-packaging. In *ACM Symposium on Document Engineering*, pages 129–138, 2011.
- [74] M. Kurant, A. Markopoulou, and P. Thiran. On the bias of breadth first search (bfs) and of other graph sampling techniques. In *Proc. 22nd International Teletraffic Congress*, pages 1–8, 2010.
- [75] N. Kushmerick. *Wrapper induction for information extraction*. PhD thesis, University of Washington, 1997. Chairperson-Weld, Daniel S.
- [76] N. Kushmerick. Wrapper induction: efficiency and expressiveness. *Artificial Intelligence*, 118(1-2):15–68, 2000.
- [77] N. Kushmerick. Finite-state approaches to web information extraction. *Proc. of 3rd Summer Convention on Information Extraction*, pages 77–91, 2002.
- [78] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *Proc. 19th International Conference on World Wide Web*, pages 591–600, Raleigh, North Carolina, USA, 2010. ACM.
- [79] A. H. F. Laender, B. A. Ribeiro-Neto, A. S. da Silva, and J. S. Teixeira. A brief survey of web data extraction tools. *SIGMOD Rec.*, 31(2):84–93, 2002.
- [80] B. Li, Q. Yang, and X. Xue. Can movies and books collaborate? cross-domain collaborative filtering for sparsity reduction. In *Proc. International Joint Conference on Artificial Intelligence*, pages 2052–2057, Innsbruck, Austria, 2009. Morgan Kaufmann Publishers Inc.
- [81] B. Li, Q. Yang, and X. Xue. Transfer learning for collaborative filtering via a rating-matrix generative model. In *Proc. International Conference on Machine Learning*, pages 617–624, Montreal, Canada, 2009.
- [82] B. Liu. Structured data extraction: Wrapper generation. *Web Data Mining*, pages 363–423, 2011.
- [83] L. Liu, C. Pu, and W. Han. XWrap: An extensible wrapper construction system for internet information. In *Proc. 16th*

- International Conference on Data Engineering*, 2000.
- [84] C. D. Manning and H. Schütze. *Foundations of statistical natural language processing*. MIT Press, Cambridge, MA, USA, 1999.
 - [85] J. Masanès. Web archiving: issues and methods. *Web Archiving*, pages 1–53, 2006.
 - [86] A. Mathes. Folksonomies-cooperative classification and communication through shared metadata. *Computer Mediated Communication*, 47(10), 2004.
 - [87] W. May, R. Himmeröder, G. Lausen, and B. Ludäscher. A unified framework for wrapping, mediating and restructuring information from the web. In *Advances in Conceptual Modeling*.
 - [88] E. Melomed, I. Gorbach, A. Berger, and P. Bateman. *Microsoft SQL Server 2005 Analysis Services (SQL Server Series)*. Sams, Indianapolis, IN, USA, 2006.
 - [89] X. Meng, D. Hu, and C. Li. Schema-guided wrapper maintenance for web-data extraction. In *Proc. of the ACM international workshop on Web information and data management*, pages 1–8, New York, NY, USA, 2003. ACM.
 - [90] A. Misllove, M. Marcon, K. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proc. 7th SIGCOMM*, pages 29–42. ACM, 2007.
 - [91] A. E. Monge. Matching algorithm within a duplicate detection system. *IEEE Techn. Bulletin Data Engineering*, 23(4), 2000.
 - [92] R. Mooney. Relational learning of pattern-match rules for information extraction. In *Proc. of the National Conference on Artificial Intelligence (AAAI 1999)*, pages 328–334, 1999.
 - [93] I. Muslea, S. Minton, and C. Knoblock. A hierarchical approach to wrapper induction. In *Proc. of the 3rd annual conference on Autonomous Agents*, pages 190–197, New York, NY, USA, 1999. ACM.
 - [94] M. Newman. The structure and function of complex networks. *SIAM review*, pages 167–256, 2003.
 - [95] E. Oro, M. Ruffolo, and S. Staab. Sxpath - extending xpath towards spatial querying on web documents. *PVLDB*, 4(2):129–140, 2010.
 - [96] S. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
 - [97] D. Perito, C. Castelluccia, M. Kaafar, and P. Manils. How unique and traceable are usernames? In *Privacy Enhancing Technologies*, pages 1–17. Springer, 2011.
 - [98] X. H. Phan, S. Horiguchi, and T. Ho. Automated data extraction from the web with conditional models. *Int. J. Bus. Intell. Data Min.*, 1(2):194–209, 2005.
 - [99] C. Plake, T. Schiemann, M. Pankalla, J. Hakenberg, and U. Leser. Alibaba: Pubmed as a graph. *Bioinformatics*, 22(19):2444–2445, 2006.
 - [100] E. Rahm and P. Bernstein. A survey of approaches to automatic schema matching. *The VLDB Journal*, 10(4):334–350, 2001.
 - [101] E. Rahm and H. H. Do. Data cleaning: Problems and current approaches. *IEEE Bulletin on Data Engineering*, 23(4), 2000.
 - [102] D. Reis, P. Golgher, A. Silva, and A. Laender. Automatic web news extraction using tree edit distance. In *Proc. of the International Conference on World Wide Web (WWW 2004)*, pages 502–511, New York, USA, 2004. ACM.
 - [103] D. Romero and J. Kleinberg. The Directed Closure Process in Hybrid Social-Information Networks, with an Analysis of Link Formation on Twitter. In *Proc. 4th International Conference on Weblogs and Social Media*, 2010.
 - [104] A. Sahuguet and F. Azavant. Building light-weight wrappers for legacy web data-sources using w4f. In *Proc. 25th International Conference on Very Large Data Bases*, pages 738–741, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
 - [105] S. Sarawagi. Information extraction. *Found. Trends databases*, 1(3):261–377, 2008.
 - [106] R. Schifanella, A. Barrat, C. Cattuto, B. Markines, and F. Menczer. Folks in folksonomies: social link prediction from shared metadata. In *Proc. 3rd ACM international conference on Web Search and Data Mining*, pages 271–280, New York, USA, 2010. ACM Press.
 - [107] S. Selkow. The tree-to-tree editing problem. *Information processing letters*, 6(6):184–186, 1977.
 - [108] Y. Shi, M. Larson, and A. Hanjalic. Tags as bridges between domains: Improving recommendation with tag-induced cross-domain collaborative filtering. In *Proc. International Conference on User Modeling, Adaption and Personalization*, Lecture Notes in Computer Science, pages 305–316, Girona, Spain, 2011. Springer.
 - [109] S. Soderland. Learning information extraction rules for semi-structured and free text. *Machine learning*, 34(1):233–272, 1999.
 - [110] S. Spiegel, J. Kunegis, and F. Li. Hydra: a hybrid recommender system [cross-linked rating and content information]. In *Proc. ACM International Workshop on Complex Networks Meet Information & Knowledge Management*, pages 75–80, Hong Kong, China, 2009.
 - [111] A. Stewart, E. Diaz-Aviles, W. Nejdl, L. B. Marinho, A. Nanopoulos, and L. Schmidt-Thieme. Cross-tagging for personalized open social networking. In *Proc. ACM International Conference on Hypertext and Hypermedia*, pages 271–278, Torino, Italy, 2009.
 - [112] M. Szomszor, I. Cantador, and H. Alani. Correlating user profiles from multiple folksonomies. In *Proc. ACM International Conference on Hypertext and Hypermedia*, pages 33–42, Pittsburgh, PA, USA, 2008. ACM.
 - [113] M. Szomszor, C. Cattuto, H. Alani, K. O’Hara, A. Baldassarri, V. Loreto, and V. Servedio. Folksonomies, the semantic web, and movie recommendation. In *Proc. European Semantic Web Conference*, Innsbruck, Austria, 2007.
 - [114] J. Turmo, A. Ageno, and N. Català. Adaptive information extraction. *ACM Comput. Surv.*, 38(2):4, 2006.
 - [115] J. Vosecky, D. Hong, and V. Shen. User identification across multiple social networks. In *Proc. International Conference on Networked Digital Technologies*, pages 360–365, Ostrava, Czech Republic, 2009. IEEE.

- [116] J. Wang, B. Shapiro, D. Shasha, K. Zhang, and K. Currey. An algorithm for finding the largest approximately common substructures of two trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):889–895, 1998.
- [117] P. Wang, W. Hawk, and C. Tenopir. Users’ interaction with world wide web resources: an exploratory study using a holistic approach. *Inf. Process. Manage.*, 36:229–251, January 2000.
- [118] G. Weikum. Harvesting, searching, and ranking knowledge on the web: invited talk. In *Proc. 2nd ACM International Conference on Web Search and Data Mining*, pages 3–4, New York, NY, USA, 2009. ACM.
- [119] C. Wilson, B. Boe, A. Sala, K. Puttaswamy, and B. Zhao. User interactions in social networks and their implications. In *Proc. 4th European Conference on Computer Systems*, pages 205–218. ACM, 2009.
- [120] W. Winkler. The state of record linkage and current research problems. In *Statistical Research Division, US Census Bureau*, 1999.
- [121] T. Winograd. *Understanding Natural Language*. Academic Press, Inc., Orlando, FL, USA, 1972.
- [122] W. Yang. Identifying syntactic differences between two programs. *Software - Practice and Experience*, 21(7):739–755, 1991.
- [123] S. Ye, J. Lang, and F. Wu. Crawling Online Social Graphs. In *Proc. 12th International Asia-Pacific Web Conference*, pages 236–242. IEEE, 2010.
- [124] R. Zafarani and H. Liu. Connecting corresponding identities across communities. In *Proc. 3rd International Conference on Weblogs and Social Media*, pages 354–357, San Jose, California, USA, 2009. The AAAI Press.
- [125] A. Zanasi. Competitive intelligence through data mining public sources. In *Competitive intelligence review*, volume 9, pages 44–54. Wiley, New York, NY, ETATS-UNIS (1990-2001) (Revue), 1998.
- [126] Y. Zhai and B. Liu. Web data extraction based on partial tree alignment. In *Proc. of the International Conference on World Wide Web (WWW 2005)*, pages 76–85, NY, 2005. ACM.
- [127] Y. Zhai and B. Liu. Structured data extraction from the web based on partial tree alignment. *IEEE Transactions on Knowledge and Data Engineering*, 18(12):1614–1628, 2006.
- [128] K. Zhang, R. Statman, and D. Shasha. On the editing distance between unordered labeled trees. *Information processing letters*, 42(3):133–139, 1992.
- [129] Y. Zhang, B. Cao, and D. Yeung. Multi-domain collaborative filtering. In *Proc. 26th Conference on Uncertainty in Artificial Intelligence*, pages 725–732, Catalina Island, California, USA, 2010. AUAI Press.
- [130] Z. Zhang, C. Zhang, Z. Lin, and B. Xiao. Blog extraction with template-independent wrapper. In *Proc. International Conference on Network Infrastructure and Digital Content*, pages 313–317. IEEE, 2010.
- [131] H. Zhao. *Automatic wrapper generation for the extraction of search result records from search engines*. PhD thesis, State University of New York at Binghamton, Binghamton, NY, USA, 2007. Adviser-Meng, Weiyi.