

Discovering Direct and Indirect Matches for Schema Elements*

Li Xu and David W. Embley
Department of Computer Science
Brigham Young University
Provo, Utah 84602, U.S.A.
{lx, embley}@cs.byu.edu

Abstract

Automating schema matching is challenging. Previous approaches (e.g. [MBR01, DDH01]) to automating schema matching focus on computing direct element matches between two schemas. Schemas, however, rarely match directly. Thus, to complete the task of schema matching, we must also compute indirect element matches. In this paper, we present a framework for generating direct as well as many indirect element matches between a source schema and a target schema. Recognizing expected data values associated with schema elements and applying schema-structure heuristics are the key ideas to computing indirect matches. Experiments we have conducted over several real-world application domains show encouraging results, yielding over 90% precision and recall for both direct and indirect element matches.

Keyword: Schema matching, data integration, schema integration, data exchange.

1 Introduction

In this paper, we focus on the long-standing and challenging problem of automating schema matching [MBR01]. Schema matching is a key operation for many applications including data integration, schema integration, message mapping in E-commerce, and semantic query processing [RB01]. Schema matching takes two schemas as input and produces a semantic correspondence between the schema elements in the two input schemas [RB01]. In this paper, we assume that we wish to map schema elements from a *source* schema into a *target* schema. In its simplest form, the semantic correspondence is a set of *direct element matches* each of which binds a source schema element to a target schema element if the two schema elements are semantically equivalent. To date, most research [BCV99, DDH01, EJJ01, LC00, MBR01, MZ98, PTU00] has focused on computing direct element matches. Such simplicity, however, is rarely sufficient, and researchers have thus proposed the use of queries over source schemas to form virtual schema elements to bind with target schema elements [BE03, MHH00]. In this more complicated form, the semantic correspondence is a set of

*This material is based upon work supported by the National Science Foundation under grant IIS-0083127.

indirect element matches each of which binds a virtual source schema element to a target schema element through appropriate *manipulation operations* over a source schema.

We assume that all source and target schemas are described using rooted conceptual-model graphs (a conceptual generalization of XML). Element nodes either have associated data values or associated object identifiers, which we respectively call *value schema elements* and *object schema elements*. We augment schemas with a variety of ontological information. For this paper the augmentations we discuss are WordNet [Mil95], sample data, and regular-expression recognizers. For each application, we construct a lightweight domain ontology [ECJ⁺99], which declares the regular-expression recognizers. We use the regular-expression recognizers to discover both direct and indirect matches between two arbitrary schemas. Based on the graph structure and these augmentations, we exploit a broad set of techniques together to settle direct and indirect element matches between a source schema and a target schema. As will be seen, regular-expression recognition and schema structure are the key ways to detect indirect element matches.

In this paper, we offer the following contributions: (1) a way to discover many indirect semantic correspondences between a source schema S and a target schema T as well as the direct correspondences and (2) experimental results of our implementation to show that it performs as well (indeed better) than other approaches for direct element matches and also performs exceptional well for the indirect matches with which we work. We present the details of our contribution as follows. Section 2 explains what we mean by direct and indirect matches between S and T . Section 3 describes a set of basic matching techniques to find potential element matches between elements in S and elements in T , and to provide confidence measures between 0 (lowest confidence) and 1 (highest confidence) for each potential match. Section 4 presents an algorithm to settle direct and indirect matches between S and T . Section 5 gives experimental results for a data set used in [DDH01] to demonstrate the success of our approach. In Section 6 we review related work, and in Section 7 we summarize, consider future work, and draw conclusions.

2 Source-to-Target Mappings

We represent all source and target schemas using rooted conceptual-model graphs. Nodes of the graph denote object and value schema elements, and edges of the graph denote relationships among object and value schema elements. The root node is a designated object of primary interest.

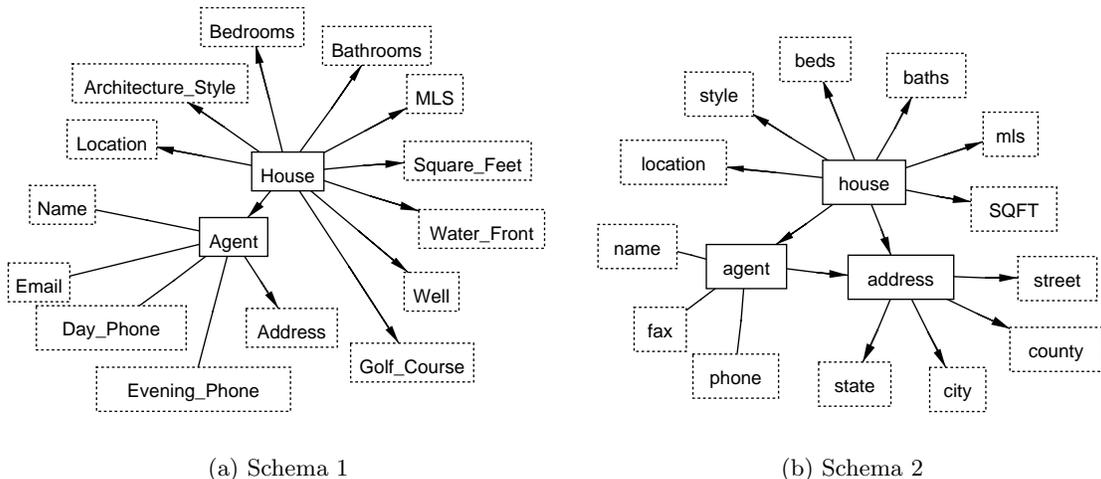


Figure 1: Schema Graphs for Schema 1 and Schema 2

Figure 1, for example, shows two schema graphs, each partially describing a real-estate application. In a schema graph we denote value schema elements as dotted boxes, object schema elements as solid boxes, functional relationships as lines with an arrow from domain to range, and nonfunctional relationships as lines without arrowheads.

The output of schema matching is a set of element mappings that match actual or virtual source schema elements with fixed target schema elements. Our source-to-target mappings allow for a variety of source derived data, including missing generalizations and specializations, merged and split values, and transformation of attributes with Boolean indicators into values.

We say that a match (s, t) is *direct* when a source schema element s and a target schema element t denote the same set of values or objects. To detect direct matches, researchers typically look for synonym matches between names of schema elements. Sometimes, however, the identification of synonyms is not enough. For example, *Location* in Figure 1(a) is the location address of a house to be sold, and *location* in Figure 1(b) is the lot description for a listed property (e.g. beach-front property). Our approach considers both schema information and data instances to help settle direct element matches, and thus largely avoids this problem of being misled by polysemy.

Although a source may not have a schema element that directly matches a target element, target facts may nevertheless be derivable from source facts. We call these correspondences *indirect* matches. When trying to detect indirect matches, we consider the following problems, which we illustrate using the schemas in Figure 1.

1. *Generalization and Specialization.* Two elements, *Day_Phone* and *Evening_Phone* in Figure 1(a) are both specializations of *phone* values in Figure 1(b). Thus, if Figure 1(b) is the target, we need the union of *Day_Phone* and *Evening_Phone*, and if Figure 1(a) is the target, we should find a way to separate the day phones from the evening phones.
2. *Merged and Split Values.* Four elements, *street*, *county*, *city*, and *state* are separate in Figure 1(b) and merged as *Location* of a house or *Address* of an agent in Figure 1(a). Thus, we need to split the values if Figure 1(b) is the target and merge the values if Figure 1(a) is the target.
3. *Schema Element Name as Value.* In Figure 1(a), the features *Water_Front* and *Golf_Course* are schema element names rather than values. The Boolean values “Yes” and “No” associated with them are not the values but indicate whether the values *Water_Front* and *Golf_Course* should be included as description values for *location* in Figure 1(b).

Currently, we use five operations over source schemas to resolve these problems.

1. *Selection.* The data values of a target schema element are a subset of the values of a source schema element.
2. *Union.* The data values of a target schema element are a superset of the values of a source schema element (usually several source schema elements). *Union* is the inverse of *Selection*.
3. *Composition.* The values of a target schema element match a concatenation of values from two or more source schema elements.
4. *Decomposition.* The values of target schema elements match a decomposition of values of a source schema element. *Decomposition* is the inverse of *Composition*.
5. *Boolean.* Attribute names with Boolean values (e.g. “Yes/No”) of a source (target) schema are values in a target (source) schema.

The recognition and specification of these operations depend on the matching techniques we describe in Sections 3 and 4. Generating operations for *Merged* and *Split Values* and for *Subsets* and *Supersets* is straightforward if we can recognize the types of matches required. For *Schema Element*

Name as Value, the resolution depends on being able to recognize the element name as a potential target value, or element values as potential target element names. Then, in harmony with the source values (e.g. “Yes”/“No”) and target element names or source element names and target values (e.g. “Yes”/“No”), we can determine the mapping.

3 Matching Techniques

In this section we explain our four basic techniques for matching: (1) terminological relationships (e.g. synonyms and hypernyms), (2) data-value characteristics (e.g. string lengths and alphanumeric ratios), (3) domain-specific, regular-expression matches (i.e. the appearance of expected strings), and (4) structure (e.g. structural similarities). For the first two techniques we obtain vectors of measures for the features of interest and then apply machine learning over these feature vectors to generate a decision rule and a measure of confidence for each generated decision. We use C4.5 [Qui93] as our decision-rule and confidence-measure generator.

3.1 Terminological Relationships

The meaning of element names provides a clue about which elements match. To match element names, we use WordNet [Fel98, Mil95] which organizes English words into synonym and hypernym sets. Other researchers have also suggested using WordNet to match attributes (e.g. [BCV99, CA99]), but have given few, if any, details.

Initially we investigated the possibility of using 27 available features of WordNet in an attempt to match a token A appearing in the name of a source schema element s with a token B appearing in the name of an target schema element t . The C4.5-generated decision tree, however, was not intuitive.¹ We therefore introduced some bias by selecting only those features we believed would contribute to a human’s decision to declare a potential attribute match, namely (f0) same word (1 if $A = B$ and 0 otherwise), (f1) synonym (1 if “yes” and 0 if “no”), (f2) sum of the distances of A and B to a common hypernym (“is kind of”) root (if A and B have no common hypernym root, the distance is defined as a maximum number in the algorithm), (f3) the number of different common hypernym roots of A and B , and (f4) the sum of the number of senses of A and B . For our training data we used 222 positive and 227 negative A - B pairs selected from attribute names

¹An advantage of decision-tree learners over other machine learning (such as neural nets) is that they generate results whose reasonableness can be validated by a human.

```

f3 <= 0: NO (222.0/26.0)
f3 > 0
| f2 <= 2: YES (181.0/3.0)
| f2 > 2
| | f4 <= 11
| | | f2 <= 5: YES (15.0/5.0)
| | | f2 > 5: NO (14.0/6.0)
| | f4 > 11: NO (17.0/2.0)

```

Figure 2: Generated WordNet Rule

found in database schemas, which were readily available to us, along with synonym names found in dictionaries. Figure 2 shows the resulting decision tree. Surprisingly, neither f_0 (same word) nor f_1 (synonym) became part of the decision rule. Feature f_3 dominates—when WordNet cannot find a common hypernym root, the words are not related. After f_3 , f_2 makes the most difference—if two words are closely related to the same hypernym root, they are a good potential match. (Note that f_2 covers f_0 and f_1 because both identical words and direct synonyms have zero distance to a common root; this helps mitigate the surprise about f_0 and f_1 .) Lastly, if the number of senses is too high ($f_4 > 11$), a pair of words tends to match almost randomly; thus the C4.5-generated rule rejects these pairs and accepts fewer senses only if pairs are reasonably close ($f_2 \leq 5$) to a common root.

The parenthetical numbers (x/y) following “YES” and “NO” for a decision-tree leaf L give the total number of training instances x classified for L and the number of incorrect training instances y classified for L . Based on the trained decision rule in Figure 2, we compute a confidence value, denoted $conf_1(s, t)$, where s is a source schema element and t is a target schema element. However, we want the feature f_0 (same word) to dominate the others and assign a perfect confidence value (1.0) for two tokens if f_0 holds. When schema element names are abbreviations, we expand them so that WordNet can recognize them. If the names of both s and t are single-word tokens, the computation of $conf_1(s, t)$ is straightforward based on the decision rule when f_0 does not hold. For a “YES” leaf L , we compute confidence factors by the formula $(x-y)/x$ where x is the total number of training instances classified for L and y is the number of incorrect training instances classified for L . For a “NO” leaf, the confidence factor is $1-(x-y)/x$, which converts “NO’s” into “YES’s” with inverted confidence values. If a schema element name is a phrase instead of a single-word token,

we select nouns from the phrase. Then if either s or t has a name consisting of multiple noun tokens, we use an injective greedy noun-match algorithm to locate the potential matching tokens between the name phrases of s and t . We compute $conf_1(s, t)$ as the average of the confidence values collected from the potential matching tokens obtained from the injective greedy algorithm.

Assuming Schema 1 in Figure 1(a) is a target schema, and Schema 2 in Figure 1(b) is a source schema, when we apply the test for terminological relationships of schema element names, the confidence value $conf_1(s, t)$ is high for the matches such as (*house*, *House*), (*beds*, *Bedrooms*), (*baths*, *Bathrooms*), (*phone*, *Day_Phone*), and (*phone*, *Evening_Phone*), as it should be. Also, the confidence of (*location*, *Location*) is high, even though the meaning is entirely different; but, as we shall see, other techniques can sort out this anomaly.

3.2 Data-Value Characteristics

Whether two sets of data have similar value characteristics provides another a clue about which elements match. Previous work in [LC00] shows that this technique can successfully help match elements by considering such characteristics as string-lengths and alphabetic/non-alphabetic ratios of alphanumeric data and means and variances of numerical data. We use features similar to those in [LC00], but generate a C4.5 decision rule rather than a neural-net decision rule. Based on the decision rule, which turns out to be lengthy but has a form similar to the decision tree in Figure 2, we generate a confidence value, denoted $conf_2(s, t)$, for each element pair (s, t) of value schema elements.

Testing the decision rule using data values associated with Schema 2 in Figure 1(b) as a source schema and Schema 1 in Figure 1(a) as a target schema, the confidence value $conf_2(s, t)$ is high for the matches such as (*beds*, *Bedrooms*), (*baths*, *Bathrooms*), (*phone*, *Day_Phone*), and (*fax*, *Day_Phone*) as expected. However, *mls* in the source and *Location* in the target tend to look alike according to the value characteristics measured, a surprise which needs other techniques to find the difference. Interestingly, the lot features in *location* of the source schema and the house locations in *Location* of the target schema do not have similar value characteristics; this is because their alphabetic/non-alphabetic ratios are vastly different, as they should be.

3.3 Expected Data Values

Whether expected values appear in a set of data provides yet another clue about which elements match. For a specific application, we can specify a lightweight domain ontology [ECJ⁺99], which includes a set of concepts and relationships among the concepts, and associates with each concept a set of regular expressions that matches values and keywords expected to appear for the concept. Then, using techniques described in [ECJ⁺99], we can extract values from sets of data associated with source and target value elements and categorize their data-value patterns based on the regular expressions declared for application concepts. The derived data-value patterns and the declared relationship sets among concepts in the domain ontology can help discover both direct and indirect matches for schema elements.

We declare the concepts and relationship sets in our lightweight domain ontologies independently of any target and source schemas. We call them lightweight for two reasons. (1) The construction of concepts and relationships is not the same as the construction of a conceptual schema in global-as-view approaches [Ull97] for integrating heterogeneous information sources. A global-as-view information-integration system maintains a global schema, and the system needs to update the global schema when new information sources become available. Thus, the global-as-view approach requires that the global schema should be complete in the sense that it embodies all the contents in the underlying information sources. We neither require nor expect that the knowledge declared in an application domain ontology is complete for the application. Moreover, (2) the objective of the regular expressions declaring expected values for application concepts is to discover corresponding concepts, not to extract items of interest [ECJ⁺99]. Since the domain ontology need not be as complete nor as exact as the declarations for a data-extraction ontology, we see our domain ontologies as being lightweight.

Figure 3 shows three components in our real-estate domain ontology, which we used to automate matching of the two schemas in Figure 1 and also for matching real-world schemas in the real-estate domain in general. The three components include an address component specifying *Address* as potentially consisting of *State*, *City*, *County*, and *Street*;² a phone component specifying *Phone* as a possible superset of *Day Phone*, *Evening Phone*, *Home Phone*, *Office Phone*, and *Cell Phone*;³

²Filled-in (black) triangles denote aggregation (“part-of” relationships).

³Open (white) triangles denote generalization/specialization (“ISA” supersets and subsets).

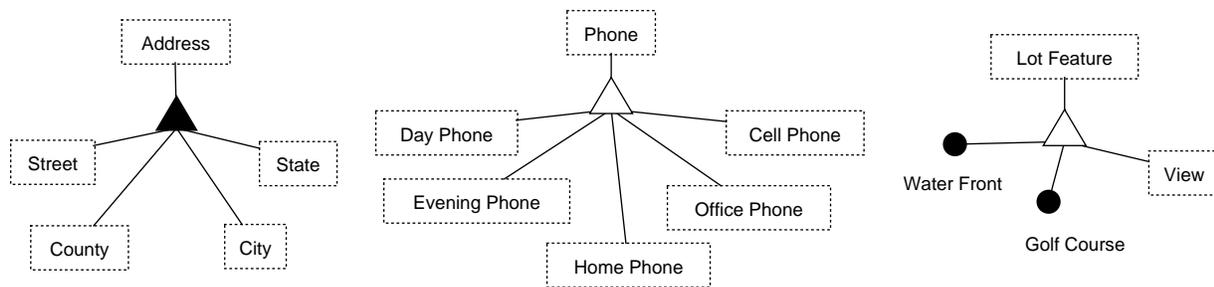


Figure 3: Application Domain Ontology (Partial)

and a lot-feature component specifying *Lot Feature* as a possible superset of *View* values and individual values *Water Front* and *Golf Course*.⁴ Behind a dotted box (or individual value), a regular-expression recognizer [ECJ⁺99] describes the expected data values for a potential application concept. The ontology explicitly declares that (1) the expected values in *Address* match with a concatenation of the expected values for *Street*, *County*, *City* and *State*; (2) the set of values associated with *Phone* is a superset of the values in *Day Phone*, *Evening Phone*, *Home Phone*, *Office Phone*, and *Cell Phone*; and (3) the set of values associated with *Lot Feature* is a superset of the values associated with the set of *View* values and the singleton-sets *Water Front* and *Golf Course*.

Provided with the domain ontology just described and a set of data values in value elements in Schema 1 in Figure 1(a) and Schema 2 in Figure 1(b), we can discover indirect matches as follows. (We first explain the idea with examples and then more formally explain how this works in general.)

1. *Composition* and *Decomposition*. Based on the *Address* declared in the ontology in Figure 3, the recognition-of-expected-values technique [ECJ⁺99] can help detect that (1) the values of both *Address* and *Location* in Schema 1 match with the ontology concept *Address*, and (2) the values of *street*, *county*, *city*, and *state* in Schema 2 match with the ontology concepts *Street*, *County*, *City*, and *State* respectively. Thus, if Schema 1 is the source and Schema 2 is the target, we can use *Decomposition* over *Address* and *Location* in the source to indirectly match with *street*, *county*, *city*, and *state* in the target. If we switch and let Schema 2 be the source and Schema 1 be the target, based on the same information, we can identify the

⁴Large black dots denote individual objects or values.

same set of indirect matching element pairs except that the manipulation operation becomes *Composition*.

2. *Union* and *Selection*. Based on the specification of the regular expression matched for *Phone*, the schema elements *Day_Phone* and *Evening_Phone* in Schema 1 match with the concepts *Day Phone* and *Evening Phone* respectively, and *phone* in Schema 2 also matches with the concept *Phone*. *Phone* in the ontology explicitly declares that the set of expected values of *Phone* is a superset of the expected values of *Day Phone* and *Evening Phone*. Thus, we are able to identify the indirect matching schema elements between *phone* in Schema 2 and *Day_Phone* and *Evening_Phone* in Schema 1. If Schema 1 is the source and Schema 2 is the target, we can apply a *Union* operation over Schema 1 to derive a virtual schema element *Phone'*, which can directly match with *phone* in Schema 2. If Schema 2 is the source and Schema 1 is the target, we may be able to recognize keywords such as *day-time*, *day*, *work phone*, *evening*, and *home* associated with each listed phone in the source. If so, we can use a *Selection* operation to sort out which phones belong in which specialization (if not, a human expert may not be able to sort these out either).
3. *Schema Element Name as Value*. Because regular-expression recognizers can recognize schema element names as well as values, the recognizer for *Lot Feature* recognizes names such as *Water_Front* and *Golf_Course* in Schema 1 as values. Moreover, the recognizer for *Lot Feature* can also recognize data values associated with *location* in Schema 2 such as “*Mountain View*”, “*City Overlook*”, and “*Water-Front Property*”. Thus, when Schema 2 is the source and Schema 1 is the target, whenever we match a target-schema-element name with a source *location* value, we can declare “Yes” as the value for the matching target concept. If, on the other hand, Schema 1 is the source and Schema 2 is the target, we can declare that the schema element name should be a value for *location* for each “Yes” associated with the matching source element.

We now more formally describe these three types of indirect matches. Let c_i be an application concept, such as *Street*, and consider a concatenation of concepts such as *Address* components. Suppose the regular expression for concept c_i matches the first part of a value v for a value schema element and the regular expression for concept c_j matches the last part of v , then we say that the

concatenation $c_i \circ c_j$ matches v . In general, we may have a set of concatenated concepts C_s match a source element s and a set of concatenated concepts C_t match a target element t . For each concept in C_s or in C_t , we have an associated hit ratio. Hit ratios give the percentage of s or t values that match (or are included in at least some match) with the values of the concepts in C_s or C_t respectively. We also have a hit ratio r_s associated with C_s that gives the percentage of s values that match the concatenation of concepts in C_s , and a hit ratio r_t associated with C_t that gives the percentage of t values that match the concatenation of concepts in C_t . To obtain hit ratios for Boolean fields recognized as schema-element names, we distribute the schema-element names over all the Boolean fields.

We decide if s matches with t directly or indirectly by comparing C_s and C_t . If C_s equals C_t , we declare a *direct* match (s, t) . Otherwise, if $C_s \supset C_t$ ($C_s \subset C_t$), we derive an *indirect* match (s, t) through a *Decomposition* (*Composition*) operation. If both C_s and C_t contain one individual concept c_s and c_t respectively, and if the values of concept c_s (c_t) are declared as a subset of the values of concept c_t (c_s), we derive an *indirect* match (s, t) through a *Union* (*Selection*) operation. When we have schema-element names as values, distribution of the name over the Boolean value fields converts these schema elements into standard schema elements with conventional value-populated fields. Thus, no additional comparisons are needed to detect direct and indirect matches when schema-element names are values. We must, however, remember the Boolean conversion for both source and target schemas to correctly derive indirect matches.

We compute the confidence value for a mapping (s, t) , which we denoted as $conf_3(s, t)$, as follows. If we can declare a direct match or derive an indirect match through manipulating *Union*, *Selection*, *Composition*, *Decomposition*, and *Boolean* for (s, t) , and the hit ratios r_s and r_t are above an accepted threshold, we output the highest confidence value 1.0 for $conf_3(s, t)$. Otherwise, we construct two vectors v_s and v_t whose coefficients are hit ratios associated with concepts in C_s and C_t . To take the partial similarity between v_s and v_t into account, we calculate a VSM [BYRN99] cosine measure $\cos(v_s, v_t)$ between v_s and v_t , and let $conf_3(s, t)$ be $(\cos(v_s, v_t) \times (r_s + r_t)/2)$.

3.4 Structure

We consider structure matching as one more technique that provides a clue about which elements match. Given the confidence measures output from the other matching techniques as a guide,

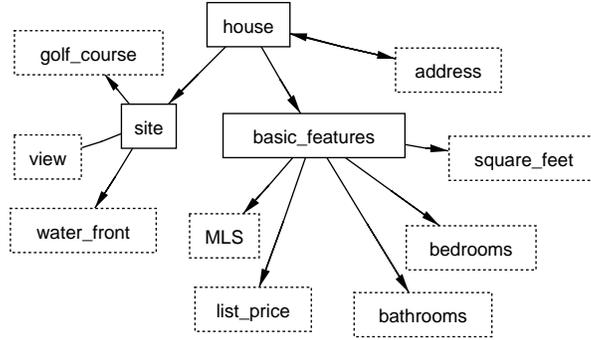


Figure 4: Schema Graph for Schema 3

structure matching determines element matches by considering contexts around schema elements. We base our structure contexts on three intuitive ideas, which we illustrate using Schemas 1 and 2 in Figure 1 and Schema 3 in Figure 4, which also partially describes a real-estate application, .

1. *Object schema element.* Two object schema elements match if their element names are similar and the value elements around them describe similar data in the two schemas. An object schema element has only object identifiers in a target or source schema. The object identifiers themselves do not describe the objects in the object element. Instead, the values of adjacent value schema elements around the object schema element describe object elements. In Figure 4, Schema 3 uses *site* to represent the lot description for a house. The confidence value $conf_1(site, address)$ computed based on terminological relationships between the two element names declares that *site* of Schema 3 is similar to *address* of Schema 2 in Figure 1(b). The data associated with the adjacent *golf_course*, *water_front*, and *view* elements together, however, represent objects about lot descriptions. In contrast, data associated with the adjacent *state*, *city*, *county*, and *street* elements together represent objects about house locations and agent contact addresses. By taking the adjacent value elements into account, *site* of Schema 3 does not match with *address* of Schema 2.
2. *Value schema element.* Adjacent object schema elements supply additional constraints for value schema elements. In Figure 1(a), *Location* denotes house locations, and *Address* denotes agent contact addresses. Even though both value schema elements are addresses, the object elements *House* and *Agent* further specialize the addresses. In Figure 4, *address* denotes house locations. Based on terminological relationships, value characterizes, and expected

data values, *address* of Schema 3 could match with either *Location* or *Address* in Schema 1. By taking the schema structure features into account, however, the value element *address* of Schema 3 should only match with the value element *Location* of Schema 1.

3. *Context extensions.* There may be functional dependencies among object elements. When an object element e functionally determines another object element e' , either directly or indirectly, we include the value elements of e' as value elements of e . Because of the functional relationship, the value elements of e' are likely to have a close relationship with e —likely much closer than value elements when there is no functional dependency. In Figure 4, the constraints in Schema 3 imply that *house* functionally determines the objects in *basic_features*. Thus, we extend context of *house* to include *MLS*, *square_feet*, *list_price*, *bathrooms*, and *bedrooms*. These elements further help specialize the data that describes a house. We use these elements in the same way we use adjacent elements.

As an example of how structure uses contexts of schema elements to help resolve schema matching, and especially how it helps identify indirect element matches, consider *address* in Schema 2 (Figure 1(b)), which contains address objects that are functionally dependent on the object schema elements *house* and *agent*. In Schema 1 (Figure 1(a)), there are two kinds of addresses: *Location*, which contains house location addresses, and *Address*, which contains agent contact addresses. Assume that Schema 2 is the source and Schema 1 is the target. By considering the value elements, we observe that *street*, *county*, *city* and *state* in Schema 2 match with both *Location* and *Address* in Schema 1 indirectly through the *Composition* operation with a confidence factor, *conf*₃. Based on this observation and on structural observations, we can declare two sets of indirect element matches. One set includes (*state*, *Location*), (*county*, *Location*), (*city*, *Location*), and (*street*, *Location*). The other set includes (*state*, *Address*), (*county*, *Address*), (*city*, *Address*), and (*street*, *Address*). For each matching element pair, we add a *Selection* operation, based on the structure, in conjunction with the *Composition* operation to separate the concatenation of *street*, *county*, *city*, and *state* in Schema 2 to match correctly with *Location* and *Address* in Schema 1.

4 Matching Algorithm

We have implemented an algorithm using our matching techniques that produces both direct and indirect matches between a source schema S and a target schema T . Figure 5 gives the algorithm, which we informally explain as follows.

Step 1: *Compute conf measures between S and T .* For each pair of schema elements (s, t) , which are either both value elements or both object elements, the algorithm computes a confidence value, $conf(s, t)$, to combine the output confidence values of the three nonstructural matching techniques.

We compute $conf(s, t)$ using the following formula.

$$conf(s, t) = \begin{cases} conf_1(s, t) , & \text{if } s \text{ and } t \text{ are object schema elements} \\ 1.0 , & \text{if } conf_3(s, t) = 1.0 \text{ and } s \text{ and } t \text{ are value schema elements} \\ w_s(conf_1(s, t)) + w_v(conf_2(s, t) + conf_3(s, t))/2 , & \text{otherwise} \end{cases}$$

In this formula, w_s and w_v are experimentally determined weights.⁵ When the confidence value $conf_3(s, t) = 1.0$, we let $conf_3$ dominate and assign $conf(s, t)$ as 1.0 and keep the detected manipulation operations (*Selection, Union, Composition, Decomposition, Boolean*) for indirect element matches. The motivation for letting $conf_3(s, t)$ dominate is that when expected values appear in both source and target schema elements and they both match well with the values we expect, this is a strong indication that the elements should match (either directly or indirectly). Since the domain ontology is not guaranteed to be complete (and may even have some inaccuracies) for a particular application domain, the confidence values obtained from the other techniques can complement and compensate for the inadequacies of the domain knowledge. This motivates the third part of the computation for $conf(s, t)$.

Step 2: *Settle object element matches.* When comparing two object element s and t , we take three factors into account: (1) the combined confidence measure $conf(s, t)$, (2) an importance similarity measure $sim_{importance}(s, t)$, and (3) a vicinity similarity measure $sim_{vicinity}(s, t)$. We can declare a matching pair (s, t) if $conf(s, t)$, $sim_{importance}(s, t)$, and $sim_{vicinity}(s, t)$ are high.

The latter two measures together represent the similarity between the contexts of s and t . We let $atoms_{direct}(e)$ denote the set of value elements directly connected to an object schema element e and

⁵The two parameters w_s , which weights schema element names, and w_v , which weights schema element values, are application dependent. Using a heuristic guide, however, we can determine the two parameters based on schemas and available data even without experimental evidence. If the schema element names are informative and the data is not self descriptive, we assign w_s as 0.8 and w_v as 0.2. On the other hand, if the schema element names are not informative and the data is semantically rich, we assign w_s as 0.2 and w_v as 0.8. For all other cases, we assign both w_s and w_v as 0.5.

Input: source schema S and target schema T
Output: a set of element matches with manipulation operations

Step 1: Compute $conf$ measures between S and T
collect the object elements in S into S_1 , and collect the value elements in S into S_2
collect the object elements in T into T_1 , and collect the value elements in T into T_2
for each (s, t) in $(S_1 \times T_1) \cup (S_2 \times T_2)$
 compute $conf_1(s, t)$ based on terminological relationships
for each (s, t) in $S_2 \times T_2$
 compute $conf_2(s, t)$ based on data-value characteristics
 compute $conf_3(s, t)$ based on expected data values
for each (s, t) in $S_1 \times T_1$
 $conf(s, t) = conf_1(s, t)$
for each (s, t) in $S_2 \times T_2$
 if $conf_3(s, t) = 1.0$ then $conf(s, t) = conf_3(s, t)$
 else
 $c_s(s, t) = conf_1(s, t)$
 $c_v(s, t) = (conf_2(s, t) + conf_3(s, t))/2$
 $conf(s, t) = c_s(s, t) \times w_s + c_v(s, t) \times w_v$

Step 2: Settle object element matches
for each s in S_1 and each t in T_1
 collect $atoms_{direct}(s)$, $atoms(s)$, $atoms_{direct}(t)$ and $atoms(t)$
for each (s, t) in $S_1 \times T_1$
 compute $sim_{vicinity}(s, t)$ and $sim_{importance}(s, t)$
 if $sim_{vicinity}(s, t) > th_{vicinity}$ and $sim_{importance}(s, t) > th_{importance}$
 and $conf(s, t) > th_{conf}$ then
 mark (s, t) as selected, mark s in S_1 , and mark t in T_1
Adjust $atoms_{direct}$ sets in S and T as follows
 for each unmarked s in S_1
 if $max_{t_i \in T_1}(sim_{vicinity}(s, t_i)) > th_{vicinity}$ then
 adjust every $atoms_{direct}(s') = atoms_{direct}(s') \cup atoms_{direct}(s)$
 where s' is a parent object schema element of s on which s is functionally dependent
 for each unmarked t in T_1
 if $max_{s_i \in S_1}(sim_{vicinity}(s_i, t)) > th_{vicinity}$ then
 adjust every $atoms_{direct}(t') = atoms_{direct}(t') \cup atoms_{direct}(t)$
 where t' is a parent object schema element of t on which t is functionally dependent
assign appropriate operations with object element matches

Step 3: Settle value element matches
for each selected (s, t) , which is a settled object element match
 for each (s', t') in $atoms_{direct}(s) \times atoms_{direct}(t)$
 if $conf(s', t') = 1.0$ then
 mark settled element match (s', t')
 mark s' and t' in $atoms_{direct}(s)$ and $atoms_{direct}(t)$ respectively
combine $conf$ measures into a single $conf$ matrix M for each pair (s'', t'') ,
 where $s'' \in atoms_{direct}(s)$ and s'' is not marked, and $t'' \in atoms_{direct}(t)$ and t'' is not marked
while there is an unsettled $conf$ measure in M greater than th_{conf}
 find the largest unsettled $conf$ measure in M
 settle $conf$ by setting it to 1, and mark $conf$ as being settled
 for each unsettled $conf'$ in the rows and columns of $conf$
 settle $conf'$ by setting it to 0, and mark $conf'$ as being settled
mark settled element matches based on the settled $conf$ measures
assign appropriate operations with value element matches

Step 4: Output element matches with manipulation operations

Figure 5: Matching Algorithm

let $atoms(e) = \bigcup_{e' \in E'} atoms_{direct}(e')$ denote the value elements of e , where E' is an object schema element set including e and other object schema elements that are functional dependent on e . We denote $atoms_{value}(S)$ and $atoms_{value}(T)$ as the sets of all value elements collected from S and T respectively. Given an experimentally determined threshold, th_{conf} ,⁶ we calculate $sim_{importance}(s, t)$ and $sim_{vicinity}(s, t)$ based on the following formulas.

$$sim_{vicinity}(s, t) = \max\left(\frac{|\{x \in atoms(s) \wedge \exists y \in atoms(t) (conf(x, y) > th_{conf})\}|}{|atoms(s)|}, \frac{|\{x \in atoms(t) \wedge \exists y \in atoms(s) (conf(y, x) > th_{conf})\}|}{|atoms(t)|} \right)$$

$$sim_{importance}(s, t) = 1.0 - \left| \frac{atoms(s)}{atoms_{value}(S)} - \frac{atoms(t)}{atoms_{value}(T)} \right|$$

Intuitively, $sim_{vicinity}$ measures the similarity of the vicinity surrounding s and the vicinity surrounding t , and $sim_{importance}$ measures the similarity of the “importance” of s and the “importance” of t where we measure the “importance” of an object node N by counting the number of value nodes related to N and all other object nodes in the functional closure of N . When the number of schema elements is largely different, it is difficult to decide the vicinity similarity based on one singular measure, $sim_{vicinity}$ [MBR01]. The conceptual analysis techniques discussed in [CAFP98] motivated $sim_{importance}$, which helps measure the context similarity from an additional perspective.

Step 3: Settle value element matches. For each matching pair (s, t) of object elements settled in Step 2, we first settle value element matches of children of s and t (or children of functionally dependent object elements of children of s and t) that match with high confidence ($conf = 1.0$). For all remaining unsettled value schema elements of s and t , we find a best possible match so long as the confidence of the match is above the threshold, th_{conf} . For each of the matches, given the structure information and the expected-value matches, we determine the appropriate operation (or sequence of operations) required to transform source schema elements into virtual elements that directly match with target schema elements.

Step 4: Output both direct and indirect element matches with manipulation operations.

⁶For any application, the computed confidence values tend to converge to a specific high measure for element matches between two schemas. Thus, we use a universal threshold value. Experimentally, we have determined that 0.7 works well across all applications.

5 Experimental Results

We evaluate the performance of our approach based on three measures: precision, recall and the F-measure, a standard measure for recall and precision together [BYRN99]. Given (1) the number of direct and indirect matches N determined by a human expert, (2) the number of correct direct and indirect matches C selected by our process described in this paper and (3) the number of incorrect matches I selected by our process, we compute the recall ratio as $R = C/N$, the precision ratio as $P = C/(C + I)$, and the F-measure as $F = 2/(1/R + 1/P)$. We report all these values as percentages.

We tested the approach proposed here using the running example in our paper and also on several real-world schemas in three different application domains. In our experiments, we evaluated the contribution of different techniques and different combinations of techniques. We always used both structure and terminological relationships because given any two schemas, these techniques always apply even when no data is available. Thus, we tested our approach with four runs on each source-target pair. In the first run, we considered only terminological relationships and structure. In the second run, we added data-value characteristics. In the third run, we replaced data-value characteristics with expected data values, and in the fourth run we used all techniques together.

5.1 Running Example

We applied the matching algorithm explained in Section 4 to the schemas in Figure 1 populated (by hand) with actual data we found in some real-estate sites on the Web. First we let Schema 2 in Figure 1(b) be the source and Schema 1 in Figure 1(a) be the target. Then, we reversed the schemas and let Schema 1 be the source and Schema 2 be the target.

Table 1 shows a summary of the results for each run in the first test where we let Schema 2 be the source and Schema 1 be the target. In the first run for the first test, the algorithm discovered all 8 direct matches correctly, but it also misclassified the source schema element *location* (meaning “view” or “on the water front” or “by a golf course”) by matching it with the target schema element *Location* (meaning address). In the first run, the algorithm also successfully discovered 2 of the 12 indirect matches—(*phone*, *Day_Phone*) and (*phone*, *Evening_Phone*)—and correctly output the *Selection* operation.⁷ In the second run, by adding the analysis of data-value

⁷We observe, however, that although we correctly generated a *Selection* operator to specialize the *phone* value

Run Nr.	Number of Matches	Number Correct	Number Incorrect	Recall %	Precision %	F-Measure %
1 (WS)	20	10	1	50%	91%	65%
2 (WCS)	20	10	0	50%	100%	67%
3 (WES)	20	20	0	100%	100%	100%
4 (WCES)	20	20	0	100%	100%	100%

W = Terminological Relationships using WordNet
C = Data-Value Characteristics
E = Expected Data Values
S = Structure

Table 1: Results for Running Example

characteristics, the false positive (*location*, *Location*) disappeared, but the algorithm generated no more indirect matches than in the first run. In both the third and fourth runs, the algorithm successfully discovered all direct and indirect matches.⁸ Especially noteworthy, we observed that our approach correctly discovered context-dependent indirect matches (e.g. (*city*, *Address*), (*state*, *Address*), ...) and appropriately produced operations composed of a combination of *Composition* and *Selection*.⁹ The result of the second test on our running example, in which we switched the schemas and let Schema 1 be the source schema and Schema 2 be the target schema, gave the same results as in Table 1.

5.2 Real-World Examples

We considered three real-world applications: *Course Schedule*, *Faculty*, and *Real Estate* to evaluate our approach. We used a data set downloaded from the LSD homepage [DDH01] for these three applications, and we faithfully translated the schemas from DTDs used by LSD to rooted conceptual-model graphs. For testing these real-world applications, we decided to let any one of the schema graphs for an application be the target and let any other schema graph for the same application be the source. Because our tests are nearly symmetrical, however, we decided not to

in Schema 2, the value transformation for *Selection* depends on keywords such as *day – time*, *day*, *work phone*, *evening*, and *home* associated with listed phone numbers. If the keywords are not available, the *Selection* operation fails to sort out the *phone* values.

⁸We correctly generated a *Selection* operator to select the right subsets of *location* (meaning “view,” etc.) in Schema 2 for *Water_Front* and *Golf_Course*, and discarded the remaining values, which were inapplicable for Schema 1. The *Selection* operator sorted out values based on the expected data values specified in the lightweight domain ontologies.

⁹The *Selection* operator sorted out the addresses composed from *state*, *city*, *county*, and *street* based on the two relationship sets *house – address* and *agent – address* in Schema 2.

Application	Number of Matches	Number Correct	Number Incorrect	Recall %	Precision %	F-Measure %
Course Schedule	128	119	1	93%	99%	96%
Faculty	140	140	0	100%	100%	100%
Real Estate	245	229	22	93%	91%	92%
All Applications	513	488	23	95%	95%	95%

Table 2: Results for Real-World Examples

test any target-source pair also as a source-target pair (as we did in our running example). We also decided not to test any single schema as both a target and a source. Since for each application there were five schemas, we tested each application 10 times. All together we tested 30 target-source pairs. For each target-source pair, we made four runs, the same four (*WS*, *WCS*, *WES*, and *WCES*) we made for our running example. All together we processed 120 runs.

Table 2 shows as summary of the results for the real-world data using all four techniques together. In two of the three applications, *Course Schedule* and *Faculty*, there were no indirect matches. For all four runs on *Faculty* every measure (recall, precision, F-measure) was 100%. For *Course Schedule*, the first and second run achieved above 90% and below 95% on all measures; and the third and fourth run gave the results for *Course Schedule* as Table 2 shows.

The *Real Estate* application exhibited several indirect matches. The problem of *Merged/Split Values* appeared twice, the problem of *Subsets/Supersets* appeared 24 times, and the problem of *Schema Element Name as Value* appeared 5 times. The experiments showed that the application of expected data values in the third and fourth run greatly affected the performance. In the first run, the measures were only about 75%. In the second run, the use of data-value characteristics improved the performance, but only a little because the measures were still below 80%. By applying expected data values in the last two runs, however, the performance improved dramatically. In the third run, the F-measures reached 91% and reached 92% by using all four techniques as Table 2 shows.

Our process successfully found all the indirect matches related to the problems of *Merged/Split Values* and *Schema Element Name as Value*. For the problem of *Subsets/Supersets*, our process correctly found 22 of the 24 indirect matches and declared two extra indirect matches.¹⁰ Over

¹⁰Of these four, three of them were ambiguous, making it nearly impossible for a human to decide, let alone a machine. In two cases there were various kinds of phones for firms, agents, contacts, and phones with and without message features, and in another case there were various kinds of descriptions and comments about a house written in

all the indirect element mappings, the three measures (recall, precision, and F-measure) were (coincidentally) all 94%.

5.3 Discussion

The experimental results show that the combination of terminological relationships and structure alone can produce fairly reasonable results, but by adding our technique of using expected data values, the results are dramatically better. Unexpectedly, the technique of using data-value characteristics did not help very much for these application domains. Our analysis of data-value characteristics is similar to the analysis in SEMINT [LC00], which produced good results for their test data. The data instances in the real-world applications we used, however, do not appear to be as regular as might be expected. The statistics are highly variant, for example, in applications such as *Course Schedule* and *Real Estate*. For these applications, a large amount of training data would be needed to train a universal decision tree required for this approach.

Some element matches failed in our approach partly because they are potentially ambiguous, and our assertions about what should and should not match are partly subjective.¹¹ Even though we tested our approach using the same test data set as in LSD [DDH01], the answer keys were generated separately and may not be the same. Furthermore, neither the experimental methodologies nor the performance measures used are the same. With this understanding, we remark that they reported approximate accuracies of 70% for *Course Schedule*, 90% for *Faculty*, 70% and 80% for the two experiments they ran on the *Real Estate* application. Thus, although our raw performance numbers are an improvement over [DDH01], we do not try to draw any final conclusion.

One obvious limitation of our approach is the need to construct an application-specific domain ontology. Currently, we manually construct these domain ontologies. As we explained in Section 3, however, these domain ontologies are lightweight and are relatively easy to construct and need not be complete. It is possible, however, to make use of statistical learning techniques to collect a set of informative and representative keywords for application concepts. Thus, without human interaction, except for some labeling, we can make use of many keywords taken from the data of the application itself and thus specify regular-expression recognizers for the application concepts

free-form text. The one clear incorrect match happened when our process unioned office and cell phones together and mapped them to phones for a firm instead of just mapping office phones to firm phones and discarding cell phones, which had no match at all in the other schema.

¹¹It is not always easy to do ground-truthing [HKL⁺01].

at least in a semi-automatic way. Furthermore, many values, such as dates, times, and currency amounts are common across many application domains and can easily be shared. Since domain ontologies appear to play an important role in indirect matching, finding ways to semi-automatically generate them is a goal worthy of some additional work.

6 Related Work

[RB01] provides a survey of several schema mapping systems. We do not repeat this work here, but instead describe work related to our approach from two perspectives: (1) work on discovering direct matches for schema elements and (2) work on discovering indirect matches for schema elements.

Direct Matches. Most of the approaches [BCV99, DDH01, EJX01, LC00, MBR01, MZ98, PTU00] to automating schema matching focus only on generating direct matches for schema elements.

- In some of our previous work [EJX01], we experimented with using data instances to help identify direct element matches. In this paper, we refine this work by adding a structural component and also extend it to the harder problem of discovering indirect matches.
- Like our approach, the LSD system [DDH01] applies a meta-learning strategy to compose several base matchers, which consider either data instances, or schema information. LSD largely exploits machine learning techniques. There are two phases in the LSD system: one is training and the other is testing. In the training phase, LSD requires training data for each matching element between two schemas for base matchers and the meta matcher. For each different application, however, both base and meta learners have to be supervised, and the supervisor must supply and mark training data to train the learners. Our approach differs in the three ways. (1) We applied machine learning algorithms only to terminological relationships and data-value characteristics. (2) Our system learned a universal decision tree for all application domains based on a domain-independent training set. Thus our system avoids the work of collecting and labeling training data for each application as in LSD. (3) To combine techniques, we let structure features guide the matching based on the results from multiple kinds of independent matches.
- SEMINT [LC00] applies neural-network learning to automating schema matching based on

instance contents. It is an element-level schema matcher because it only considers attribute matching without taking the structure of schemas into account.

- The structure matching algorithm in Cupid [MBR01] motivated our structure matching algorithm. Cupid, however, does not properly handle two schemas that are largely different. Moreover, the structure matching algorithm Cupid matches two schemas using a bottom-up strategy. Our matching algorithm discovers direct and indirect matches using a top-down strategy.
- ARTEMIS [BCV99], DIKE [PTU00], and Cupid [MBR01] exploit auxiliary information such as synonym dictionaries, thesauri, and glossaries. All their auxiliary information is schema-level—does not consider data instances. In our approach, the auxiliary information including data instances and domain ontologies provide a more precise characterization of the actual contents of schema elements. The imported dictionary we use, WordNet, is readily available and no work is required to produce thesauri as in other approaches.

Indirect Matches. Some work on indirect matches is starting to appear [BE03, MBR01, MHH00, MWJ99], but researchers are only beginning to scratch the surface of the multitude of problems.

- Both Cupid [MBR01] and SKAT [MWJ99] can generate global 1 : n indirect matches [RB01]. To illustrate what this means, if in Figure 1 we let Schema 1 be the source and Schema 2 be the target, and if we make *address* a value element rather than an object element and discard *street*, *county*, *city*, and *state* in Schema 2, Cupid can match both *Address* and *Location* in the source directly with the modified *address* in the target. Thus Cupid can generate a global 1 : n indirect match through a *Union* operation. Our approach, however, can find indirect matches for *Location* and *Address* in the source with *street*, *county*, *city*, and *state* in the target based on finding expected data values and using the *Decomposition* operator as well as the *Union* operator, something which is not considered in Cupid.
- The Clio system [MHH00] introduces an interactive mapping creation paradigm based on value correspondence that shows how a value of a target schema element can be created from a set of values of source elements. A user or DBA, however, is responsible to manually input the value correspondences.

- [BE03] proposes a mapping generator to derive an injective target-to-source mapping including indirect matches in the context of information integration. The mapping generator raises specific issues for a user’s consideration. The mapping generator, however, has not been implemented. Our work therefore builds on and is complimentary to the work in [BE03].

7 Conclusion

We presented a framework for automatically discovering both direct matches and many indirect matches between sets of source and target schema elements. In our framework, multiple techniques each contribute in a combined way to produce a final set of matches. Techniques considered include terminological relationships, data-value characteristics, expected values, and structural characteristics. We detected indirect element matches for *Selection*, *Union*, *Composition*, and *Decomposition* operations as well as *Boolean* conversions for *Schema-Element Names as Values*. We base these operations and conversions mainly on expected values and structural characteristics. Additional indirect matches, such as arithmetic computations and value transformations, are for future work. We also plan to semi-automatically construct domain ontologies used for expected values, automate application-dependent parameter tuning, and test our approach in a broader set of real-world applications. As always, there is more work to do, but the results of our approach for both direct and indirect matching are encouraging, yielding over 90% in both recall and precision.

References

- [BCV99] S. Bergamaschi, S. Castano, and M. Vincini. Semantic integration of semistructured and structured data sources. *SIGMOD Record*, 28(1):54–59, March 1999.
- [BE03] J. Biskup and D.W. Embley. Extracting information from heterogeneous information sources using ontologically specified target views. *Information Systems*, 28(1), 2003. To appear, currently a manuscript version of this paper is at <http://www.deg.byu.edu/papers/int.pdf>.
- [BYRN99] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, Menlo Park, California, 1999.
- [CA99] S. Castano and V. De Antonellis. ARTEMIS: Analysis and reconciliation tool environment for multiple information sources. In *Proceedings of the Convegno Nazionale Sistemi di Basi di Dati Evolute (SEBD’99)*, pages 341–356, Como, Italy, June 1999.

- [CAFP98] S. Castano, V. De Antonellis, M.G. Fugini, and B Pernici. Conceptual schema analysis: Techniques and applications. *ACM Transactions on Database Systems*, 23(3):286–333, September 1998.
- [DDH01] A. Doan, P. Domingos, and A. Halevy. Reconciling schemas of disparate data sources: A machine-learning approach. In *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data (SIGMOD 2001)*, pages 509–520, Santa Barbara, California, May 2001.
- [ECJ⁺99] D.W. Embley, D.M. Campbell, Y.S. Jiang, S.W. Liddle, D.W. Lonsdale, Y.-K. Ng, and R.D. Smith. Conceptual-model-based data extraction from multiple-record Web pages. *Data & Knowledge Engineering*, 31(3):227–251, November 1999.
- [EJX01] D.W. Embley, D. Jackman, and Li Xu. Multifaceted exploitation of metadata for attribute match discovery in information integration. In *Proceedings of the International Workshop on Information Integration on the Web (WIIW’01)*, pages 110–117, Rio de Janeiro, Brazil, April 2001.
- [Fel98] C. Fellbaum. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, Massachusetts, 1998.
- [HKL⁺01] J. Hu, R. Kashi, D. Lopresti, G. Nagy, and G. Wilfong. Why table ground-truthing is hard. In *Proceedings of the Sixth International Conference on Document Analysis and Recognition*, pages 129–133, Seattle, Washington, September 2001.
- [LC00] W. Li and C. Clifton. SEMINT: A tool for identifying attribute correspondences in heterogeneous databases using neural networks. *Data & Knowledge Engineering*, 33(1):49–84, 2000.
- [MBR01] J. Madhavan, P.A. Bernstein, and E. Rahm. Generic schema matching with Cupid. In *Proceedings of the 27th International Conference on Very Large Data Bases (VLDB’01)*, pages 49–58, Rome, Italy, September 2001.
- [MHH00] R. Miller, L. Haas, and M.A. Hernandez. Schema mapping as query discovery. In *Proceedings of the 26th International Conference on Very Large Databases (VLDB’00)*, pages 77–88, Cairo, Egypt, September 2000.
- [Mil95] G.A. Miller. WordNet: A lexical database for English. *Communications of the ACM*, 38(11):39–41, November 1995.
- [MWJ99] P. Mitra, G. Wiederhold, and J. Jannink. Semi-automatic integration of knowledge sources. In *Proceedings of the Second International Conference on Information Fusion (FUSION 99)*, Sunnyvale, California, July 1999.
- [MZ98] T. Milo and S. Zohar. Using schema matching to simplify heterogeneous data translation. In *Proceedings of the 24th International Conference on Very Large Data Bases (VLDB-98)*, pages 122–133, New York City, New York, August 1998.
- [PTU00] L. Palopoli, G. Teracina, and D. Ursino. The system DIKE: Towards the semi-automatic synthesis of cooperative information systems and data warehouses. In *Proceedings of ADBIS-DASFAA 2000*, pages 108–117, 2000.

- [Qui93] J.R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, California, 1993.
- [RB01] E. Rahm and P.A. Bernstein. A survey of approaches to automatic schema matching. *The VLDB Journal*, 10(4):334–350, 2001.
- [Ull97] Jeffrey D. Ullman. Information integration using logical views. In Foto N. Afrati and Phokion Kolaitis, editors, *Proceedings of the 6th International Conference on Database Theory (ICDT'97)*, volume 1186 of *Lecture Notes in Computer Science*, pages 19–40, Delphi, Greece, January 1997. Springer-Verlag.