



MIT Sloan School of Management

**Working Paper 4411-01
November 2001**

DIRECT: A System for Mining Data Value Conversion Rules from Disparate Data Sources

Weiguo Fan, Hongjun Lu, Stuart E. Madnick and David Cheung

© 2001 by Weiguo Fan, Hongjun Lu, Stuart E. Madnick and David Cheung.
All rights reserved. Short sections of text, not to exceed two paragraphs,
may be quoted without explicit permission, provided that
full credit including © notice is given to the source.

This paper also can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:

http://ssrn.com/abstract_id=377900

DIRECT: A System for Mining Data Value Conversion Rules from Disparate Data Sources

Weiguo Fan

Dept. of Computer & Information Systems
Univ. of Michigan Business School
wfan@umich.edu

Hongjun Lu

Dept. of Computer Science
Hong Kong Univ. of Science and Technology
luhj@cs.ust.hk

Stuart E. Madnick

Sloan School of Management
MIT
smadnick@mit.edu

David Cheung

Dept. of Computer Science
Hong Kong Univ.
dcheung@cs.hku.hk

Abstract

The successful integration of data from autonomous and heterogeneous systems calls for the resolution of semantic conflicts that may be present. Such conflicts are often reflected by discrepancies in attribute values of the same data object. In this paper, we describe a recently developed prototype system, **DIRECT** (*DI*scovering and *RE*conciling *C*onflic*T*s). The system mines data value conversion rules in the process of integrating business data from multiple sources. The system architecture and functional modules are described. The process of discovering conversion rules from sales data of a trading company is presented as an illustrative example.

Keywords: data integration, data mining, semantic conflicts, data visualization, statistical analysis, data value conversion

1 Introduction

Data mining or knowledge discovery from databases (KDD) is defined by Fayyad et al. [11] as the non-trivial process of identifying valid, novel, potentially useful and ultimately understandable hidden structure, patterns, knowledge from data. Organizations are taking advantage of “data mining” techniques to leverage the vast amount of data to make better business decisions [26]. For example, Data mining has been used for customer profiling in CRM and customer service support [17], credit card application approval, fraud detection, telecommunication network monitoring, market-basket analysis [11], healthcare quality assurance [36], and many other decision making areas [3].

In this paper, we are extending the data-mining paradigm to a new application area: data integration from disparate sources [10, 21], in which database administrators often find it very hard to make decisions on the integration of those semantically related attributes [28].

It is known that the successful integration of data from autonomous and heterogeneous systems calls for the resolution of semantic conflicts that may be present. Such conflicts are often reflected by discrepancies in attribute values of the same data object. Traditional approaches to achieving semantic interoperability among heterogeneous and autonomous systems can be identified as either *tightly-coupled* or *loosely-coupled* [34]. In tightly-coupled systems, semantic conflicts are identified and reconciled a priori in one or more shared schema, against which all user queries are formulated. This is typically accomplished by the system integrator or DBA who is responsible for the integration project. In a loosely-coupled system, conflict detection and resolution is the responsibility of users, who must construct queries that take into account potential conflicts and identify operations needed to circumvent them. In general, conflict detection and reconciliation is known to be a difficult and tedious process since the semantics of data are usually present implicitly and often ambiguous. This problem poses even greater difficulty when the number of sources increases exponentially and when semantics of data in underlying sources changes over time. To provide support for information exchange data among heterogeneous and autonomous systems, one approach is to use context associated with the data [13, 4]. The basic idea is to partition *data sources* (databases, web-sites, or data feeds) into groups, each of which shares the same *context*, the latter allowing the semantics of data in each group to be codified in the form of *meta-data*. A *context mediator* takes on the responsibility for detecting and reconciling semantic conflicts that arise when information

is exchanged between systems (data sources or receivers). This is accomplished through the comparison of the contexts associated with sources and receivers engaged in data exchange [5].

As attractive as the above approach may seem, it requires data sources and receivers participating in information exchange to provide the required contexts. In the absence of other alternatives, this codification will have to be performed by the system integrator. This, however, is an arduous process, since data semantics is not always explicitly documented, and often evolves with the passage of time. This motivated our work reported in this paper: discovering and reconciling conflicts in disparate information sources by mining the data themselves. The rationale is simple. Most semantic conflicts, such as different data types and formats, units, granularity, synonyms, different coding schemes, etc., exist because data from different sources are stored and manipulated under different context, defined by the systems and applications. Therefore, those conflicts should be uniformly present in the data. That is, for a real world entity, although the values for the same attribute may be different since they are obtained from different sources, such differences should be present similarly in all entities. It would be reasonable to expect that such discrepancies can be identified and the relationships between conflict attributes can be discovered.

This paper reports the results of our work on discovering and reconciling semantic conflicts using data mining techniques to help DBAs to make better decisions on resolving data value related conflicts. Our proposed framework is implemented in a prototype system called **DI-RECT**. Our focus for the moment is on business and financial data, though the framework can be easily extended for other kinds of data. There are a number of reasons to start with business data. First, such data are easily available. Second, the relationships among the attribute values for business data are relatively simple compared to engineering or scientific data. Most of such relationships are linear or product of attributes. On the other hand, business data do have some complex factors. For example, monetary figures are often rounded up to the unit of currencies. Such rounding errors are in fact random noises that are mixed with the value conflicts caused by different context. While the basic techniques implemented in the prototype system are known, our contribution is to integrate various techniques (with necessary modifications and extensions) to provide a solution to a practical problem which has been thought difficult to solve.

1.1 Related work

Our work is closely related to work on discovering and resolving data value conflicts for data integration in the context of heterogeneous and autonomous database systems. In an earlier work, Dayal proposed the use of aggregate functions, e.g. average, maximum, minimum, etc. to resolve discrepancies in attribute values [8]. DeMichiel proposed to use virtual attributes and partial values to solve the problem of failing to map an attribute value to a definite value [9]. However, no details were given on how to map conflicting values in to the common domain of virtual attributes. Tseng, Chen, Yang further generalized the concept of partial values into probabilistic partial values to capture the uncertainty in attributes values [37]: The possible values of an attribute are listed and probabilities are given to indicate their likelihood. Lim, Srisvastava and Shekhar proposed an extended relational model based on *Dempster-Shafer Theory of Evidence* to deal with the situation where uncertain information arises when the database integration process requires information not directly represented in the component databases but can be obtained through some of the data [22]. The extended relation uses evidence sets to represent uncertainty information, which allows probabilities to be attached to subsets of possible domain values. Scheuermann and Chong adopted a different view of conflicting attribute values: different values mean different roles the attribute is performing [32, 33]. Therefore, it is not necessary to reconcile the conflicts. What is needed is to be able to use appropriate values for different applications. The work of Agarwal *et. al* addressed the same problem addressed in this paper: resolving conflicts in non-key attributes [2]. They proposed an extended relational model, flexible relation to handle the conflicting data values from multiple sources. No attempts were made to resolve the conflicts by converting the values.

As evident from the above discussion, most of the work in the existing literature have placed their emphasis on determining the value of an attribute involving semantic conflicts, while considerably less work on discovering and resolving conflicts through the use of conversion functions has been reported. Sciore, Siegel and Rosenthal have proposed the use of conversion functions to solve the semantic heterogeneity problem [35]. However, the details of discovering such functions were not provided.

1.2 Our proposed data mining approach

Our approach is distinct from previous efforts in the study of semantic conflicts in various ways. First, we begin with a simpler classification scheme of conflicts presented in data values. We classify such conflicts into two categories: *context dependent* and *context independent*. While context dependent conflicts are a result of disparate interpretations inherent in different systems and applications, context independent conflicts are often caused by errors and external factors. As such, context dependent conflicts constitute the more predictable of the two and can often be rectified by identifying appropriate *data value conversion rules*, or simply *conversion rules*, which we will shortly define. These rules can be used for data exchange or integration directly.

Second, our primary contribution in this paper is the development of a methodology and associated techniques that “mine” data conversion rules from data originating from disparate systems that are to be integrated. The approach requires a training data set consisting of tuples merged from data sources to be integrated or exchanged. Each tuple in the data set should represent a real world entity and its attributes (from multiple data sources) model the properties of the entity. If semantic conflicts exist among the data sources, the values of those attributes that model the same property of the entity will have different values. A mining process first identifies attribute sets each of which involves some conflicts. After identifying the relevant attributes, models for possible conversion functions, the core part of conversion rules are selected. The selected models are then used to analyze the data to generate candidate conversion functions. Finally, a set of most appropriate functions are selected and used to form the conversion rules for the involved data sources.

The techniques used in each of the above steps can vary depending on the property of data to be integrated. In this paper, we describe a statistics-based prototype system for integrating financial and business data, *DIRECT* (*DI*scovering and *RE*conciling *C*onflic*T*s). The system uses partial correlation analysis to identify relevant attributes, Bayesian information criterion for candidate model selection, and robust regression for conversion function generation. Conversion function selection is based on the support of rules. Experiment conducted using both synthetic data and a real world data set indicated that the system successfully discovered the conversion rules among data sources containing both context dependent and independent conflicts.

The contributions of our work can be summarized as follows. First, our classification scheme for semantic conflicts is more practical than previous proposals. For those context dependent conflicts, proposed data value conversion rules can effectively represent the quantitative relationships among the conflicts. Such rules, once defined or discovered, can be used in resolving the conflicts during data integration. Second, a general approach for mining the data conversion rules from actual data values is proposed. A prototype system has been developed to fully automate the whole discovery process. While the aim of the system is to provide automatic discovery, It can also be used in an interactive way to allow the user to fine-tune the system and speed up the whole mining process. Experiment results on both synthetic data and real world data set demonstrated the promising of our new proposed approach.

The remainder of the paper is organized as follows. We first define data value conflicts conversion rules in Section 2. Section 3 describes the architecture and the functional modules of **DIRECT**. Section 4 describes the techniques implemented in various modules. Section 5 presents two example processes of mining conversion rules from both synthetic data and a real trading company data. Section 5 concludes the paper.

2 Data value conflicts and conversion rules

In this section, we define data value conflicts and conversion rules. We will use the relational model for ease of explanation without losing generality. Furthermore, we assume that the entity identification problem, i.e., identifying the records representing the same real world entity from multiple data sources is solved using other methods [38, 23]. Therefore, attributes of different sources can be merged into one relation with the attributes modelling the properties of the entity. Each tuple in the relation represents a real world entity. If semantic conflicts exist among the data sources, the values of those attributes that model the same property of an entity will be different. We call such difference in data attribute values *data value conflict*.

As an illustrative example, Table 1 lists sample tuples from two relations, *stock* and *stk_rpt*:

stock (**stock**, currency, volume, high, low, close);
stk_rpt (**stock**, price, volume, value);

where relation *stk_rpt* is produced from *stock* by stock brokers for its clients based on the daily quotation of stock prices.

Insert Table 1 around here

As mentioned earlier, we assume that the entity identification problem has been largely solved. Therefore, stocks with the same code refer to the same company's stock. For example, tuple with $stock = 100$ in relation $stock$ and tuple with $stock = 100$ in relation stk_rpt refer to the same stock. If we know that both $stk_rpt.price$ and $stock.close$ are the closing price of a stock of the days, and similarly, both $stock.volume$ and $stk_rpt.volume$ are the number of shares traded during the day, it will be reasonable to expect that, for a tuple $s|s \in stock$ and a tuple $t|t \in stk_rpt$, if $s.stock = t.stock$, then $s.close = t.price$ and $s.volume = t.volume$. However, from the sample data, this does not hold. In other words, by analyzing the data, we can get the conclusion that, *data value conflicts* exist between the two data sources. A bit formally, we can define data value conflicts as follows:

Definition Given two data sources DS_1 and DS_2 and two attributes A_1 and A_2 modelling the same property of a real world entity type in DS_1 and DS_2 respectively, if $t_1 \in DS_1$ and $t_2 \in DS_2$ represent the same real world instance of the object but $t_1.A_1 \neq t_2.A_2$, then we say that a data value conflict exists between DS_1 and DS_2 .

To resolve data value conflicts, we need to not only identify such conflicts, but also discover the quantitative relationships among the attribute values involving such conflicts. For ease of explanation, we adopt the notation of Datalog rules for representing such relationship and call them *data value conversion rules* (or simply *conversion rules*) which take the form $head \leftarrow body$. The head of the rule is a predicate representing a relation in one data source. The body of a rule is a conjunction of a number of predicates, which can either be extensional relations present in underlying data sources, or built-in predicates representing arithmetic or aggregate functions. Some examples of data conversion rules are described below.

Example: For the example given above, we can have the following data value conversion rule:

$$stk_rpt(stock, price, volume, value) \leftarrow$$
$$exchange_rate(currency, rate), stock(stock, currency, = volume_in_K, high, low, close),$$
$$price = close * rate, volume = volume_in_K * 1000, value = price * volume.$$

where *exchange-rate* is another relation in the data source of stock. ¹

Example: For conflicts caused by synonyms or different representations, it is always possible to create lookup tables which form part of the conversion rules. For example, to integrate the data from two relations $D_1.student(sid, sname, major)$ and $D_2.employee(eid, ename, salary)$, into a new relation $D_1.std_emp(id, name, major, salary)$, we can have a rule

$$D_1.std_emp(id, name, major, salary) \leftarrow \\ D_1.student(id, name, major), D_2.employee(eid, name, salary), D_1.same_person(id, eid).$$

where *same_person* is a relation that defines the correspondence between the student id and employee id. Note that, when the size of the lookup table is small, it can be defined as rules without bodies. One widely cited example of conflicts among student grade points and scores can be specified by the following set of rules:

$$D_1.student(id, name, grade) \leftarrow \\ D_2.student(id, name, score), score_grade(score, grade). \\ score_grade(4, 'A'). \\ score_grade(3, 'B'). \\ score_grade(2, 'C'). \\ score_grade(1, 'D'). \\ score_grade(0, 'F').$$

One important type of conflicts mentioned in literature in business applications is *aggregation conflict* [19]. Aggregation conflicts arise when an aggregation is used in one database to identify a set of entities in another database. For example, in a transactional database, there is a relation $sales(date, customer, part, amount)$ that stores the sales of parts during the year. In an executive information system, relation $sales_summary(part, sales)$ captures total sales for each part. In some sense, we can think *sales* in relation *sales_summary* and *amount* in *sales* are attributes with the same semantics (both are the sales amount of a particular part). But there are conflicts as the *sales* in *sales_summary* is the summation of the amount in the *sales* relation. To be able to define conversion rules for attributes involving such conflicts, we can extend the traditional Datalog to include aggregate functions. In this example, we can have the following rules

¹In fact, the relation of *exchange-rate* can be discovered from the given relation *stock* and *stk-rpt*.

$$\begin{aligned} & \text{sales_summary}(\text{part}, \text{sales}) \leftarrow \\ & \text{sales}(\text{date}, \text{customer}, \text{part}, \text{amount}), \text{sales} = \text{SUM}(\text{part}, \text{amount}). \end{aligned}$$

where $\text{SUM}(\text{part}, \text{amount})$ is an aggregate function with two arguments. Attribute part is the *group-by* attribute and amount is the attribute on which the aggregate function applies. In general, an aggregate function can take $n + 1$ attributes where the last attribute is used in the aggregation and others represent the “group-by” attributes. For example, if the sales_summary is defined as the relation containing the total sales of different parts to different customers. The conversion rule could be defined as follows:

$$\begin{aligned} & \text{sales_summary}(\text{part}, \text{customer}, \text{sales}) \leftarrow \\ & \text{sales}(\text{date}, \text{customer}, \text{part}, \text{amount}), \text{sales} = \text{SUM}(\text{part}, \text{customer}, \text{amount}). \end{aligned}$$

We like to emphasize that it is not our intention to argue about appropriate notations for data value conversion rules and their respective expressive power. For different application domains, the complexity of quantitative relationships among conflicting attributes could vary dramatically and this will require conversion rules having different expressive power (and complexity). The research reported here is primarily driven by problems reported for the financial and business domains; for these types of applications, a simple Datalog-like representation has been shown to provide more than adequate representations. For other applications (e.g., in the realm of scientific data), the form of rules could be extended. Whatever the case may be, the framework set forth here can be used profitably to identify conversion rules that are used for specifying the relationships among attribute values involving conflicts.

3 DIRECT: the system architecture

Insert Figure 1 around here

In this section, we briefly describe the architecture of the system, **DIRECT** (*DI*scovering and *RE*conciling *C*onflic*T*s), developed at NUS with the objective of mining data value conversion rules. As shown in Figure 1, the system consists of a data visualization module and

a set of conversion rule discovery modules. Conversion rule discovery modules of **DIRECT** include *Relevant Attribute Analysis*, *Candidate Model Selection*, *Conversion Function Generation*, *Conversion Function Selection*, *Conversion Rule Formation* and some planned future modules, like aggregation rule discovery, etc.. In case no satisfactory conversion functions are discovered, training data is reorganized and the mining process is re-applied to confirm that there are indeed no conversion functions. The system also tries to learn from the mining process by storing used model and discovered functions in database (the *Model Base*) to assist later mining activities. Figure 2 depicts the rule discovery process.

Insert Figure 2 around here

- *Relevant attribute analysis* is the first step of the data value conversion rule mining process. For a given attribute in the integrated data set, the relevant attribute analysis module will determine the other attributes in the data set that are semantically equivalent or required for determining the values of semantically equivalent attributes. In the previous *stock* and *stk_rpt* example, *stock.close* and *stk_rpt.price* are semantically equivalent attributes because both of them represent the last trading price of the day. If we want to find the attributes that are related with *stock.close*, relevant attribute analysis module will find that it is related with *stk_rpt.price* and one latent attribute *exchange_rate*. As shown in Figure 2, meta data about the data sources may be used to help the relevance analysis. For example, data types of attributes, an easily available type of meta data, can be used to reduce the search space for semantically relevant attributes.
- *Candidate model selection* is the second module of the conversion rule mining process. To find a data value conversion rule among semantically-related attributes, various models must first be assigned to specify their quantitative relationships. Since the model space may be of a very large size,² it's essential that we can limit the model search space to make our approach viable and practical in real practice. Candidate model selection is

²For a linear model involving 10 independent attributes, the model space will be of the size $2^{10} = 1024$.

such a module to find the best potential models for the data from which the conversion functions will be generated.

- *Conversion function generation* is a module that can efficiently estimate the model parameters and goodness of those candidate models that are generated by the candidate model selection module. Note that the data set may contain a lot of outliers, or error data points, the conversion function generation module must have a very robust property against those outliers. That is, it should still give good estimations of the model parameters even in a contaminated environment.

- *Conversion function selection and Conversion rule formation*

It is often the case that the mining process generates more than one conversion functions because it is usually difficult for the system to determine the optimal one. The conversion function selection module needs to develop some measures or heuristics to select the best conversion functions from the candidates. With the selected functions, some syntactic transformations are applied to form the data conversion rules as specified in Section 2.

4 System implementation

In this section, we give a detailed description of the implementation of our **DIRECT** system. Most of techniques we used are from statistics field, with modification to meet our special needs of conversion rule mining.

4.1 Data visualization

Data visualization is one of the most important modules of the system. Before the users start the mining process, they can first use the data visualization module to visualize their data and gain some idea about the distribution of their values. As shown in Figure 3, we provide several statistics-lines on the visualization diagram: minimum, maximum and mean value lines to help users to quickly identify the required information.

Insert Figure 3 around here

Data visualization is complementary to the whole conversion rule discovery modules. A typical example is the relevant attribute analysis, which is described in a separate section. As shown in Figure 4, scatter plot, which is very powerful in examining the relationship among two attributes, is provided to help users to further determine which attributes should be included in the final model. This can help to restrain the model search space, and thus can greatly improve the system’s performance in terms of speed and accuracy.

4.2 Relevant attribute analysis

The basic techniques used to measure the (linear) association among numerical variables in statistics is *correlation analysis* and *regression*. Given two variables, X and Y and their measurements $(x_i, y_i); i = 1, 2, \dots, n$, the strength of association between them can be measured by *correlation coefficient* $r_{x,y}$.

$$r_{x,y} = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2} \sqrt{\sum(y_i - \bar{y})^2}} \quad (1)$$

where \bar{x} and \bar{y} are the mean of X and Y , respectively. The value of r is between -1 and +1 with $r = 0$ indicating the absence of any linear association between X and Y . Intuitively, larger values of r indicate a stronger association between the variables being examined. A value of r equal to -1 or 1 implies a perfect linear relation. While correlation analysis can reveal the strength of linear association, it is based on an assumption that X and Y are the only two variables under the study. If there are more than two variables, other variables may have some effects on the association between X and Y . *Partial Correlation Analysis (PCA)*, is a technique that provides us with a single measure of linear association between two variables while adjusting for the linear effects of one or more additional variables. Properly used, partial correlation analysis can uncover spurious relationships, identify intervening variables, and detect hidden relationships that are present in a data set.

It is true that correlation and semantic equivalence are two different concepts. A person’s height and weight may be highly correlated, but it is obvious that height and weight are different in their semantics. However, since semantic conflicts arise from differences in the representation schemes and these differences are uniformly applied to each entity, we expect a high correlation to exist among the values of semantically equivalent attributes. Partial correlation analysis can at least isolate the attributes that are likely to be related to one

another for further analysis.

The limitation of correlation analysis is that it can only reveal linear associations among numerical data. For non-linear associations or categorical data, other measures of correlation are available and will be integrated into the system in the near future.

A snap-shot of output of relevant attribute analysis is shown in Figure 4.

Insert Figure 4 around here

4.3 Candidate model selection

As discussed above, a model is required before performing regression. Given a set of relevant attributes, a model specifies the number of attributes that should actually be included and the basic relationship among them.

A number of techniques have been developed to automatically select and rank models from a set of attributes, like t-test, step-wise regression. All of these methods are available in the popular statistical packages, like SAS, SPSS, S-PLUS[24]. However, these model selection methods tend to favor those models that involve more attributes, which make them very vulnerable in the real modelling process [27]. We adopt the approach proposed by Raftery [27]. To assess a model, Raftery introduced the BIC, Bayesian Information Criterion, an approximate to Bayes factors used to compare two models based on Bayes's theorem. For different regression functions, the BIC takes different forms. In the case of linear regression, BIC_k of a model M_k can be computed as

$$BIC = N\log(1 - R_k^2) + p_k\log N \quad (2)$$

where N is the number of data points, R_k^2 is the value of adjusted R^2 for model M_k and p_k is the number of independent attributes. Using the BIC values of models, we can rank the models. The smaller the BIC is, the better the model is to fit the data. One important principle in comparing two nested models M_k and M_{k+1} , where k is the number of independent attributes, is so called Occam's Window. The essential idea is that, if M_k is better than M_{k+1} , model M_{k+1} is removed. However, if model M_{k+1} is better, it requires a certain "difference" between

two models to cause M_k to be removed. That is, there is an area, the Occam's Window, where M_{k+1} is better than M_k but not better enough to cause M_k being removed. The size of Occam's Window can be adjusted. Smaller Occam's Window size will cause more models to be removed. A illustrative example on synthetic data in Section 5 is provided to show its impact on the model selection process. Please refer to [27] for more detailed discussion on BIC criterion and its computation.

4.4 Conversion function generation

With a given set of relevant attributes, $\{A_1, \dots, A_k, B_1, \dots, B_l\}$ with A_1 and B_1 being semantically equivalent attributes, the conversion function to be discovered

$$A_1 = f(B_1, \dots, B_l, A_2, \dots, A_k) \quad (3)$$

should hold for all tuples in the training data set. The conversion function should also hold for other unseen data from the same sources. This problem is essentially the same as the problem of learning quantitative laws from the given data set.

Various approaches have been proposed in the literature. They can be classified into two categories: parametric approach and non-parametric approach [12]. Even though non-parametric approach has the merit that it does not require many statistical assumptions, it suffers from that fact that it is very hard to explain the relationship it discovered. Moreover, chances exist it can approximate spurious relationship even though there are no explicit function relationships existed at all [15]. Since the purpose of our data mining process is to discover the underlying conversion function that can be used for later data integration, only functions of explicit forms like linear or linear with several orders of interaction, etc. should be of our primary concern. So in the later implementation of our system, we will concentrate on using only the parametric methods, especially, parametric regression methods for the conversion function discovery purpose.

Regression analysis is one of the most well studied method for discovering such quantitative relationships among variables from a data set [7] [1]. If we view the database attributes as variables, the conversion functions we are looking for among the attributes are nothing more than regression functions. For example, the equation $stk_rpt.price = stock.close * rate$ can be viewed as a regression function, where $stk_rpt.price$ is the response (dependent) variable, $rate$

and *stock.close* are predictor (independent) variables. In general, there are two steps in using regression analysis. The first step is to define a model, which is a prototype of the function to be discovered. For example, a linear model of p independent variables can be expressed as follows:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p + \mu \quad (4)$$

where $\beta_0, \beta_1, \dots, \beta_p$ are model parameters, or *regression coefficients*, and μ is an unknown random variable that measures the departure of Y from exact dependence on the p predictor variables.

For a given model, the model parameters β_1, \dots, β_p for the regression model can be easily estimated using a variety of statistical techniques. The essence of these techniques is to fit all data points to the regression function by determining the coefficients. The goodness of the discovered function with respect to the given data is usually measured by the *coefficient of determination*, or R^2 . The value R^2 ranges between 0 and 1. A value of 1 implies perfect fit of all data points with the function. We argued earlier that the conversion function for context dependent conflicts are uniformly applied to all tuples (data points) in a data set, hence all the data points should fit to the discovered function ($R^2 = 1$). However, this will be true only if there are no noises or errors, or *outliers*, in the data. In real world, this condition can hardly be true. One common type of noise is caused by rounding of numerical numbers. This leads to two key issues in conversion function generation using regression:

- A robust regression technique that can produce the functions even when noises and errors exist in data; and
- The acceptance criteria of the regression functions generated.

We will address the first issue in this subsection. The second issue will be discussed in the next subsection.

The most frequently used regression method is Least Squares (LS) Regression. However, it is only efficient and effective for the case where there are no outliers in the data set. To deal with outliers, various robust/resistant regression methods were proposed [31], including Least Median Squares (LMS) Regression [29], Least Quartile Squares (LQS) Regression and Least Trimmed Squares (LTS) Regression [30]. We implemented the LTS regression method. LTS regression, unlike the traditional LS regression that tries to minimize the sum of squared residuals of all data points, minimizes a subset of the ordered squared residuals, and leaves

out those large squared residuals, thereby allowing the fit to stay away from those outliers and leverage points. Compared with other robust techniques, like LMS, LTS also has a very high breakdown point: 50%. That is, it can still give a good estimation of model parameters even half of the data points are contaminated. More importantly, a faster algorithm exists to estimate the parameters [6].

4.5 Conversion function selection

As mentioned in the earlier subsection, we need to define a criteria that can be used to select the best conversion function. To address this issue, we defined a new measure, *support*, to evaluate the goodness of a discovered regression function (conversion function) based on recent work of Hoeting, Raftery and Madigan [16]. A function generated from the regression analysis is accepted as a conversion function only if its support is greater than a user specified threshold γ . The *support* of a regression function is defined as

$$support = \frac{Count\{y'_i | (round(|y_i - y'_i|)/scale) \leq \epsilon\}}{N} \quad (5)$$

where $i \in (1, N)$, y_i is the actual value of the dependent attributes for the i_{th} point, y'_i is the predicted value of the dependent value using the function for the i_{th} point, ϵ is a user-defined parameter that represents the required prediction accuracy, N is the number of the total points in the data set and *Count* is a function that returns the number of data points satisfying the conditions. *scale* in Equation 5 is a robust estimate of the residuals from the regression function, defined as follows:

$$scale = 1.4826 * median(abs((Y - Y') - median(Y - Y'))) \quad (6)$$

where *median()* is a function that returns the median value of its vector. Y and Y' are the vector of the actual value and the predicted value using the regression function for the dependent attribute respectively.

The *support* indicates the percentage of data points whose residuals are within certain range of the scale estimate, $\epsilon * scale$. In other words, those points with residuals greater than $\epsilon * scale$ are identified as outliers. To determine whether a regression function should be accepted, user specifies a minimum support, γ . The system only generates those functions with *support* $>$ γ . Two parameters, ϵ and γ have different meanings. ϵ *in fact specifies*

the requirement of the conversion precision required in our case. Depending on whether the detected outliers are genuine outliers or data points still with acceptable accuracy, the value of ϵ and γ can be adjusted. The default values for *gamma* and *epsilon* are 0.80 and 4 respectively.

4.6 Conversion rule formation

We have briefly discussed the major techniques currently implemented in **DIRECT**. The last step, conversion rule formation, is merely a syntactic transformation and the details are omitted here. A snap-shot of the conversion rule formation is shown in Figure 5. If there are any potential outliers having been identified, they will also be reported by this module accordingly.

Insert Figure 5 around here

4.7 Training data reorganization

It is possible that no satisfactory conversion function is discovered with a given relevant attribute set because no such conversion function exists. However, there is another possibility from our observations: the conflict cannot be reconciled using a single function, but is reconcilable using a suitable collection of different functions. To accommodate for this possibility, our approach includes a training data reorganization module. The training data set is reorganized whenever a single suitable conversion function cannot be found. The reorganization process usually partitions the data into a number of partitions. The mining process is then applied in attempting to identify appropriate functions for each of the partition taken one at a time. This partitioning can be done in a number of different ways by using simple heuristics or complex clustering techniques, e.g. partition based on the categorical attributes, or partition based on the distribution of the dependent attributes, or even more complex, using the clustering algorithm to cluster the data set into different groups [18]. The partition should be done by the system in an automatic way, or can be offered to users as options to allow them to select. Currently we have implemented a simple heuristic that partitions the data set based on categorical attributes present in the data set. The rationale is, if multiple functions exist

for different partitions of data, data records in the same partition must have some common property, which may be reflected by values assumed by some attributes within the data being investigated.

5 Experiments

To validate the efficacy of our proposed approach, two experiments were run using different data sets. The first data set is a synthetic data set which contains no outliers, while the second may contain outliers that we don't know a priori. Both of these two experiments are done fully automatically using our **DIRECT** system without any user intervention. Some of the intermediary results are also presented for illustration purpose.

5.1 A illustrative example using synthetic data

In this subsection, we use a synthetic data set as an example to illustrate the workings of the prototype system. We use the example schema *stock* and *stk_rpt* mentioned at the beginning of the paper to generate the training data set.

stock (**scode**, currency, volume, high, low, close);
stk_rpt (**scode**, price, volume, value);

Tuples in relation *stk_rpt* are in fact derivable from relation *stock*. This is used to simulate the data integration process: The broker receives stock information and integrates it into his/her own stock report. For this experiment, we created a data set of 6000 tuples. The domain of the attributes are listed in Table 2

Insert Table 2 around here

5.1.1 Relevant attribute analysis

The zero-order correlation analysis was first applied to the training data. The results are shown in Table 3.

Insert Table 3 around here

If the correlation efficient between two attributes is greater than the threshold, which was set to 0.1 in the experiment, they were considered relevant. Therefore, three relevant attribute sets are formed as follows:

Set	Attribute	Correlated Attributes
1	A3	A6
2	A4	A6, A7, A8, A9
3	A2	A7, A8, A9

Partial correlation analysis was conducted by controlling A6, A7, A8 and A9 in turn to see whether there are more attributes that should be included in the obtained relevant attribute sets. In this example, there were no such attributes; and the relevant attribute sets remained the same.

5.1.2 Candidate model selection and conversion function generation

Each of the relevant attribute sets was used to generate conversion functions in turn. For convenience, we explain the two processes, together.

Set 1: {A3, A6}

Set 1 contains only two variables. There is only one possible model ($A3 = A6$). Only the following function was obtained with 100% support:

$$A3 = 1000 * A6. \tag{7}$$

Set 2: {A4, A6, A7, A8, A9}

A4 is correlated to four other attributes. Without any prior knowledge, and considering only linear model with first-order and second-order terms (an acceptable practice for most financial and business data), the initial model includes all the attributes.

$$A4 = A6 + A7 + A8 + A9 + A6*A7 + A6*A8 + A6*A9 + A7*A8 + A7*A9 + A8*A9 \tag{8}$$

Table 4 lists the output of the model selection module for different Occam’s Window sizes.

It can be seen that with larger window size, the number of candidate models selected increases and more attributes were included in the model. When window size is 20 and 30, only one model with attributes A_6 and A_9 was selected. When the window size increased to 40, four models were selected and attributes A_8 and A_7 appeared in some of the models.

Insert Table 4 around here

In the conversion function generation process, the selected model was used to generate conversion functions. For the models in Table 4, the system in fact did not report any functions. By further examining the process, we found that all the models resulted in functions with support lower than specified threshold, Table 5.

Insert Table 5 around here

As mentioned earlier, one possibility that a selected model does not generate any conversion function is that there exist multiple functions for the data set. To discover such multiple functions, the data set should be partitioned. In our implementation, a simple heuristic is used, that is, to partition the data using categorical attributes in the data set. After the data was partitioned based on a categorical attribute A_5 , the model selected are shown in Table 6.

Insert Table 6 around here

The conversion function generation module estimates the coefficients for each of the models selected for each partition. There is only one conversion function reported for each partition, since all the coefficients for the terms other than $A_6 * A_9$ were zero. The results are summarize in Table 7.

Insert Table 7 around here

Set 3: {A2, A7, A8, A9}

The process for relevant attribute set 3 is similar to what described for set 2. The initial model used is

$$A2 = A7 + A8 + A9 + A7 * A8 + A7 * A9 + A8 * A9 \quad (9)$$

and the model selection module selected 10 models without producing functions with enough support. Using the data sets partitioned using A5 and the conversion functions listed in Table 8 were obtained.

Insert Table 8 around here

5.1.3 Conversion function selection and data conversion rule formation

Since there is only one set of candidate functions obtained for each set of relevant attribute set with 100% support. The functions generated were selected to form the data conversion rules. By some syntactic transformation, we can obtain the following data conversion rules for our example:

```
stk_rpt(stock, price, rpt_volume, value) ←
    stock(stock, currency, stk_volume, high, low, close),
    price = rate * close, rpt_volume = 1000 * stk_volume,
    value = rate * 1000 * stk_volume * close, exchange_rate(currency, rate).
exchange_rate(0, 0.4).
exchange_rate(1, 0.7).
exchange_rate(2, 1.0).
exchange_rate(3, 1.8).
exchange_rate(4, 5.8).
```

It is obvious that the discovered rule can be used to integrate data from *stock* to *stk_rpt*.

5.2 An example of mining process using a real world data set

In this section, we describe an example mining process using **DIRECT** in which a set of real world data collected from a small trading company is used.

Insert Table 9 around here

To allow us to focus on mining data value conversion rules, a program was written to extract data from the database into a single data sets with the 10 attributes shown in Table 9. Attribute *A1* to *A8* are from a transaction database, *T*, which records the details of each invoice billed, hence all the amounts involved are in the original currency. Attribute *A9* is from a system for cost/profit analysis, *C*, which captures the sales figure in local currency exclusive of tax. Attribute *A10* is from the accounting department, *A*, where all the monetary figures are in the local currency. Therefore, although attributes *A3*, *A9* and *A10* have the same semantics, i.e., all refer to the amount billed in an invoice, their values are different. In the experiment, we tried to discover the conversion functions among those conflicting attributes.

From the context of the business, the following relationship among the attributes should exist:

$$A.amount = T.amount * T.exchange_rate. \quad (10)$$

$$C.amount = (T.amount - T.GST_amount) * T.exchange_rate \quad (11)$$

$$= A.amount - T.exchange_rate * T.GST_amount. \quad (12)$$

The goods and service tax (GST), is computed based on the amount charged for the sales of goods or services. As *C.amount* is in the local currency and all transaction data are in the original currency, we have the following relationship:

$$T.GST_amount = C.amount / T.exchange_rate * T.GST_rate \quad (13)$$

where GST rate depends on the nature of business and clients. For example, exports are exempted from GST (tax rate is 0%) and domestic sales are taxed in a fixed rate of 3% in our case.

The data set collected contains 7,898 tuples corresponding to the invoices issued during the first 6 months of a financial year. The discovering process and the results are summarized in the following subsections.

5.3 Transaction data versus analysis data

Relevant attribute analysis: As our current system only works for numeric data, the categorical attributes are not included in analysis. That is, among eight attributes from data source T , only attributes $A3$, $A5$, $A7$ and $A8$ are taken into the relevant attribute analysis. The results of the partial correlation analysis are shown in Table 10.

Insert Table 10 around here

Note that, the correlation coefficients between $A10$ and $A5$, $A10$ and $A7$ are rather small in the zero-order test. However, this alone is not sufficient to conclude that $A10$ is not related to $A5$ and $A7$. By the PCA test with controlling $A3$, the correlation between $A10$ and $A7$ became rather obvious and the correlation coefficient between $A10$ and $A5$ also increased. With one more PCA test that controls $A8$, the correlation between $A10$ and $A5$ became more obvious. Thus, all the four attributes are included into the quantitative relationship analysis.

Candidate model selection and conversion function generation: The candidate model selection and conversion function generation module produced the following function with 100% support.

$$A10 = A3 * A7 \tag{14}$$

5.4 Transaction data versus accounting data

Similarly, tests were conducted among attribute $A9$ and $A3$, $A5$, $A7$ and $A8$ to discover the relationship among $A.amount$ and $T.amount$. Table 11 listed the correlation coefficients obtained from the partial correlation analysis.

Insert Table 11 around here

Again, the following function was obtained with 100% support.

$$A9 = A3 * A7 - A7 * A8 \quad (15)$$

5.5 Tests with partitioned data

In the above results, we did not find any relationship involving attribute $A5$ which shows high correlation with both $A9$ and $A10$. Since the analysis on the whole data set did not generate any functions, we partitioned the data according to categorical attributes. One categorical attribute, $A4$ is named as *sales_type*, and is used to partition the data set first. There are two values for *sales_type*: 1, 2. The results obtained from the partitioned data are listed in Table 12.

Insert Table 12 around here

The functions with higher support in each group, i.e., Function 1 and 3 are selected as the conversion functions. We can see that they in fact represent the same function. Using the attribute numbers instead of the name, we can rewrite Equation 12 as

$$A9 = A3 * A7 - A8 * A7$$

The tax rate, $A5$, is stored as the percentage, Equation 13 can be rewritten as

$$A8 = \frac{A9}{A7} * 0.01 * A5$$

Thus we have

$$\begin{aligned} A9 &= A3 * A7 - \frac{A9}{A7} * 0.01 * A5 * A7 \\ &= A3 * A7 - 0.01 * A9 * A5 \end{aligned}$$

Therefore,

$$A9 = \frac{A3 * A7}{1 + 0.01 * A5} \quad (16)$$

From the data, we have

$$A5 = \begin{cases} 3 & \text{for } A4 = 1 \\ 0 & \text{for } A4 = 2 \end{cases}$$

Substituting this into Equation 16, we have

$$A9 = \begin{cases} \frac{A3 * A7}{1.03} = 0.9708738 * A3 * A7 & \text{for } A4 = 1 \\ A3 * A7 & \text{for } A4 = 2 \end{cases}$$

which are the functions 1 and 3. In other words, although the two functions have different forms, they do represent the original relationship among the data values specified by Equation 16.

What seems unnatural is that the support of Function 1 is not 100% as it should be. We found the following two tuples listed in Table 13 that were identified as the outliers. They are indeed the tuples with errors contained in the original data: Both tuples have incorrect value of GST amount. ³

Insert Table 13 around here

5.6 Discussions

By going through the mining process of two sample data sets using DIRECT system, we demonstrated that we could not only successfully discover the underlying conversion functions and resolve the semantic conflicts, but also address the data quality issue by successful identification of those erroneous records (outliers) existed in the databases.

There are several points we want to emphasize here on the possible usage of, or extensions to our proposed framework:

³We realized that the system used in the company allows data entry clerk to modify GST values calculated and prompted by the system. It is obvious that for these two invoices, incorrect tax values were entered.

- Robust regression is used as the primary methodology to uncover the true relationships among those semantically related attributes. By “robust”, we mean that it does NOT suffer from the situation where the error terms in the regression model do not follow a normal distribution as required in traditional least-square linear regression analysis. When there are many outliers in the data set, traditional LS regression requires users to check the normality assumptions and identify the outliers by going through a series of diagnostics steps before running the regression analysis. But even if they following these procedures, there are still some potential outliers in the data set that can not be uncovered a priori (refer to page 5-8 in [29] for detailed examples) . This is no longer the case in *LTS* robust regression since *LTS* can automatically detect and neglect those outliers in the regression analysis. Besides, it’s been shown that *LTS* robust regression is a type of non-parametric regression method which does not require very rigid linear independence and normality assumptions [29].
- At the moment, **DIRECT** supports the mining of conversion functions of a linear form with multiple order interaction terms, which is very common these days because a lot of current data warehouses and heterogeneous database systems store a large amount of business and financial transaction data. The purpose of the mining process is to mine the “original” or “meaningful” functions involving numerical attributes in the databases as we mentioned in Section 2, and to resolve the semantic conflicts among these attributes. There are many other data mining techniques proposed in the literature [14], e.g. statistical approximation using non-parametric methods like neural networks, kernel methods, etc. that could be used for this purpose. But since their results are harder to interpret and encode when we perform the data integration task, they are not implemented in our system. These techniques, however, are very useful for situations when we do not care very much about the final functional forms or when the functional form is of a highly complex and nonlinear form and is very hard to detect using conventional regression methods. The discussion of these techniques to these situations are beyond the scope of this paper.
- Data mining is an interactive process. Both of the two examples shown above are retrospective studies to validate the efficacy of our approach. The interestingness of the conversion rules we learned can be easily interpreted and verified. In real situation, the

data set could be very large. In this case, the entire data set could be sampled and splited into two sub data sets: training data and test data, following machine learning experimental design approach [25]. The training data is used to discover the possible candidate models which are presented to the user for further selection. The test data is used for final verification and selection of those candidate models. These features will be incorporated into future releases of our **DIRECT** system. Currently, the rules learned by **DIRECT** after the retrospective study are sent to the DBAs for final verification and selection to help them make better decisions.

- It is known that that there are also cases where the relationship among attributes are not linear in functional form. We are currently working on extending our system **DIRECT** to make nonlinear conversion function discovery modules using domain knowledge possible. Similar to [26], it can support the domain knowledge by knowledge representation using AI techniques. It can also support the novel function discovery using more advanced inductive discovery technique like Genetic Programming [20]. Because of the uniqueness of the nonlinear function discovery due to its very high complexity and computational intensity, the results will be reported in another paper.

6 Conclusions

In this paper, we addressed the problem of discovering and resolving data value conflicts for data integration. We first proposed a simple classification scheme for data value conflicts based on how they can be reconciled. We argue that those conflicts caused by genuine semantic heterogeneity can be reconciled systematically using data value conversion rules. A general approach for discovering data conversion rules from data automatically was proposed. A prototype system **DIRECT** was developed and tested using both synthetic and real world data set. The result is very promising. It can not only be used to discover the underlying data value conversion rules interactively or automatically, but also can be used as a tool to improve the quality of the data because of its capability to identify outliers presented in the data set.

The various modules implemented in our system are based on the approaches developed in the Statistics literature. Some of the individual modules are also implemented in the commercial statistical packages, like SPSS, SAS, S-PLUS, etc. Our system, however, is different from these systems in the following ways:

- **DIRECT** is used specifically for mining the data conversion rules. The modules are fully optimized and integrated into one system that can be run both automatically and interactively. While most of the statistical packages can not.
- Our system doesn't require the user to have many statistical background to understand and use our system. It has very nice and intuitive GUI (Graphical User Interface). All the procedures can be run by simply clicking one button. Almost all of the current statistical packages require that you have at least basic statistics knowledge to be able to operate their systems. And you need to learn on how to program using their own programming or scripting languages.
- The model selection module using BIC measure, as well as robust regression module and outlier detection module using genetic algorithm, are not available in most of the statistical packages.
- Our system supports the SQL syntax that can be used to extract the data directly from different databases. The final conversion rules can also be stored directly as meta data information into the databases.

It's not our intention to say that the methods we adopted in our implementation in each of the discovery modules are the best ones. New learning techniques keep coming up in statistics, mathematics, pattern recognition, and machine learning. However, the general discovery framework proposed in this paper, however, should be applicable to all the other function discovery processes. Future developments of **DIRECT** are the following:

- Design and implement a module that can discover nonlinear conversion functions as well as automatic power/log transformations of data.
- Design and implement a module that can discover conversion rules involving aggregate functions (like, Sum, Avg., Max. etc.) that are often presented in business and financial data.
- Currently, data repartition module only supports heuristic based on categorical attributes, i.e. partition the data using single categorical attribute or their combinations. More advanced techniques, like data classification using decision tree, neural networks,

or data clustering as discussed in [25] could be extended and implemented to support this non-trivial task.

- Data value conversion rules are now defined among data sources. It has the drawback that a large number of such rules have to be defined if we are to integrate data from a large number of sources. We are exploring some other useful schemes to improve the scalability of our system.

7 Acknowledgements

The authors wish to thank Editor-in-chief, associate editor and anonymous reviewers for their helpful and constructive suggestions for improvement of this paper.

References

- [1] A. Afifi and V. Clark, *Computer-Aided Multivariate Analysis*. Chapman & Hall, third edition (1996).
- [2] S. Agarwal, A. M. Keller, G. Wiederhold, and K. Saraswat, Flexible relation: An approach for integrating data from multiple, possibly inconsistent databases, in: *Proc. IEEE Intl Conf on Data Engineering*, Taipei, Taiwan (March 1995).
- [3] R. Brachman, T. Khabaza, W. Kloesgen, G. Piatetsky-Shapiro, and E. Simoudis, Mining business databases, *Communications of ACM*, 39(11):42–48 (1996).
- [4] S. Bressan, K. Fynn, C. Goh, M. Jakobisiak, K. Hussein, H. Kon, T. Lee, S. Madnick, T. Pena, J. Qu, A. Shum, and M. Siegel, The COntext INterchange mediator prototype, in: *Proc. ACM SIGMOD/PODS Joint Conference*, Tucson, AZ (1997).
- [5] S. Bressan, C. Goh, S. Madnick, and M. Siegel, A procedure for context mediation of queries to disparate sources, in: *Proceedings of the International Logic Programming Symposium* (October 12-17 1997).
- [6] P. Burns, A genetic algorithm for robust regression estimation, *StatScience Technical Note* (1992).
- [7] S. Chatterjee and B. Price, *Regression Analysis by Example*. John Wiley & Sons, second edition (1991).
- [8] U. Dayal, Processing queries with over generalized hierarchies in a multidatabase system, in: *Proceedings of VLDB Conference*, pages 342–353 (1983).

- [9] L. G. Demichiel, Resolving database incompatibility: an approach to performing relational operations over mismatched domains, *IEEE Trans on Knowledge and Data Engineering*, 1(4):485–493 (1989).
- [10] W. Fan, H. Lu, D. Cheung, and S. Madnick, Discovering and reconciling data value conflicts for data integration, *Information Systems (to appear)* (2001).
- [11] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth, From data mining to knowledge discovery: an overview, in: U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*. AAAI/MIT Press (1996).
- [12] J. Friedman, Multivariate adaptive regression splines, *The Annals of Statistics*, 19(1):1–141 (1991).
- [13] C. Goh, S. Bressan, S. Madnick, and M. Siegel, Context interchange: Representing and reasoning about data semantics in heterogeneous systems, *Sloan School Working Paper #3928*, Sloan School of Management, MIT, 50 Memorial Drive, Cambridge MA 02139 (Oct 1996).
- [14] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers (2000).
- [15] T. Hastie and R. Tibshirani, Nonparametric regression and classification, in: *From statistics to Neural Networks: Theory and Pattern Recognition Applications*, Springer-Verlag (1994).
- [16] J. Hoeting, A. Raftery, and D. Madigan, A method for simultaneous variable selection and outlier identification, *Technical Report 9502*, Department of Statistics, Colorado State University (1995).
- [17] S. Hui and G. Jha, Data mining for customer service support, *Information & Management*, 38(1):1–14 (2000).
- [18] A. K. Jain and R. C. Dubes, *Algorithms for Clustering Data*. Prentice Hall (1988).
- [19] V. Kashyap and A. Sheth, Schematic and semantic similarities between database objects: A context-based approach, *VLDB Journal*, 5(4) (October 1996).
- [20] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA (1992).
- [21] E.-P. Lim and R. Chiang, The integration of relationship instances from heterogeneous databases, *Decision Support Systems*, 29:153–167 (2000).
- [22] E.-P. Lim, J. Srivastava, and S. Shekhar, An evidential reasoning approach to attribute value conflict resolution in database integration, *IEEE Transactions on Knowledge and Data Engineering*, 8(5) (Oct. 1996).
- [23] E.-P. Lim, J. Srivastava, S. Shekhar, and J. Richardson, Entity identification problem in database integration, in: *Proceedings of the 9th IEEE Data Engineering Conference*, pages 294–301 (1993).

- [24] A. Miller, *Subset Selection in Regression*. Chapman and Hall (1990).
- [25] T. M. Mitchell, *Machine Learning*. McGraw Hill (1997).
- [26] B. Padmanabhan and A. Tuzhilin, Unexpectedness as a measure of interestingness in knowledge discovery, *Decision Support Systems*, 27:303–318 (1999).
- [27] A. Raftery, Bayesian model selection in social research, *Sociological Methodology*, pages 111–196 (1995).
- [28] P. Robertson, Integrating legacy systems with modern corporate applications, *Communications of ACM*, 40(5):39–46 (1999).
- [29] P. Rousseeuw, Least median of squares regression, *Journal of the American Statistical Association*, 79:871–880 (1984).
- [30] P. Rousseeuw and M. Hubert, Recent developments in progress, in: *L1-Statistical Procedures and Related Topics*, volume 31, pages 201–214, Hayward, California. Institute of Mathematical Statistics Lecture Notes-Monograph Series (1997).
- [31] P. Rousseeuw and A. Leroy, *Robust Regression and Outlier Detection*. John Wiley & Sons (1987).
- [32] P. Scheuermann and E. I. Chong, Role-based query processing in multidatabase systems, in: *Proceedings of International Conference on Extending Database Technology*, pages 95–108 (1994).
- [33] P. Scheuermann, W.-S. Li, and C. Clifton, Dynamic integration and query processing with ranked role sets, in: *Proc. First International Conference on Interoperable and Cooperative Systems (CoopIS'96)*, pages 157–166, Brussels, Belgium (June 1996).
- [34] P. Scheuermann, C. Yu, A. Elmagarmid, H. Garcia-Molina, F. Manola, D. McLeod, A. Rosenthal, and M. Templeton, Report on the workshop on heterogeneous database systems, *ACM SIGMOD RECORD*, 4(19):23–31 (Dec 1990), Held at Northwestern University, Evanston, Illinois, Dec 11-13, 1989. Sponsored by NSF.
- [35] E. Sciore, M. Siegel, and A. Rosenthal, Using semantic values to facilitate interoperability among heterogeneous information systems, *ACM Transactions on Database Systems*, 19(2):254–290 (June 1994).
- [36] M. Tsechansky, N. Pliskin, G. Rabinowitz, and A. Porath, Mining relational patterns from multiple relational tables, *Decision Support Systems*, 27:177–195 (1999).
- [37] F. S. Tseng, A. L. Chen, and W.-P. Yang, Answering heterogeneous database queries with degrees of uncertainty, *Distributed and Parallel Databases: An International Journal*, 1(3):281–302 (1993).
- [38] Y. Wang and S. Madnick, The inter-database instance identification problem in integrating autonomous systems, in: *Proceedings of the Sixth International Conference on Data Engineering* (1989).

BIOGRAPHY

Weiguo Fan is currently a Ph.D Candidate at the University of Michigan Business School. He received B.E. in Information Science and Engineering in 1995 from Xi'an Jiaotong University, Xi'an, P.R. China, and MSc degree in Computer Science in 1997 from the National University of Singapore, Singapore. His research interests include data mining, text(web) mining, information retrieval on the WWW, and information integration.

Hongjun Lu is an associate professor at the Hong Kong University of Science and Technology. He received his PhD degree in Computer Science from the University of Wisconsin, Madison. His research interests include query processing and optimization, parallel and distributed database systems, and knowledge discovery and data mining. He has published more than 80 papers in various database conferences and journals. He is currently a member of the ACM SIGMOD Advisory Committee and a trustee of the VLDB Endowment. He has served as a programme committee member for many leading database conferences like SIGMOD, VLDB, ICDE and as a reviewer for major database journals and conferences.

Stuart E. Madnick is the John Norris Maguire Professor of Information Technology and Leaders for Manufacturing Professor of Management Science at the MIT Sloan School of Management. He is also an affiliate member of the MIT Laboratory for Computer Science and a member of the Executive Committee of the MIT Center for Information Systems Research. His current research interests include connectivity among disparate distributed information systems, database technology, and software project management. He is the author or co-author of over 200 books, articles, or reports on these subjects, including the classic textbook, *Operating Systems* (McGraw-Hill), and the book, *The Dynamics of Software Development* (Prentice-Hall). He has been active in industry, making significant contributions as one of the key designers and developers of projects such as IBM's VM/370 operating system and Lockheed's DIALOG information retrieval system.

David W. Cheung received the MSc and PhD degree in computer science from Simon Fraser University, Canada, in 1985 and 1989, respectively. He also received the BSc degree in mathematics from the Chinese University of Hong Kong. From 1989 to 1993, he was with Bell Northern Research, Canada, where he was a member of the scientific staff. Currently, he is an associate professor of the Department of Computer Science at the University of Hong Kong. His research interests include distributed databases, spatial databases, data mining and data warehousing.

Table 1: Example tuples from relation *stock* and *stk_rpt*

stock						stk_rpt			
stock	currency	volume	high	low	close	stock	price	volume	value
100	1	438	100.50	91.60	93.11	100	65.18	438000	28548840.00
101	3	87	92.74	78.21	91.35	101	164.43	87000	14305410.00
102	4	338	6.22	5.22	5.48	102	31.78	338000	10741640.00
104	1	71	99.94	97.67	99.04	104	69.33	71000	4922430.00
111	0	311	85.99	70.22	77.47	111	30.99	311000	9637890.00
115	2	489	47.02	41.25	41.28	115	41.28	489000	20185920.00
120	3	370	23.89	21.09	22.14	120	39.85	370000	14744500.00
148	3	201	23.04	19.14	21.04				
149	1	113	3.70	3.02	3.32				

Table 2: Example relations

Attr. No.	Relation	Name	Value Range
A1	stock, stk_rpt	scode	[1, 500]
A2	stk_rpt	price	stock.close*exchange_rate[stock.currency]
A3	stk_rpt	volume	stock.volume * 1000
A4	stk_rpt	value	stk_rpt.price * stk_rpt.volume
A5	stock	currency	1, 2, 3, 4, 5, random
A6	stock	volume	20-500, uniform distribution
A7	stock	high	[stock.close, 1.2*stock.close]
A8	stock	low	[0.85*stock.close, stock.close]
A9	stock	close	[0.50, 100], uniform distribution

Table 3: Correlation coefficients from the zero-Order correlation analysis

	<i>A6</i>	<i>A7</i>	<i>A8</i>	<i>A9</i>
<i>A2</i>	0.0291	0.3933	0.4006	0.4050
<i>A3</i>	1.0000	-0.0403	-0.0438	-0.0434
<i>A4</i>	0.3691	0.2946	0.2994	0.3051

Table 5: Functions discovered for selected models

Model	Function generated	Support (%)
$A4 = A6 * A9$	$A4 = 708.04 * A6 * A9$	81.73
$A4 = A7 + A6 * A9$	$A4 = 9980 * A7 + 678.05 * A6 * A9$	81.22
$A4 = A8 + A6 * A9$	$A4 = 10875.13 * A8 + 685.14 * A6 * A9$	81.08
$A4 = A9 + A6 * A9$	$A4 = 9970.88 * A9 + 677.84 * A6 * A9$	81.31
$A4 = A6 * A9 + A7 * A8$	$A4 = 685.90 * A6 * A9 + 102.15 * A7 * A8$	81.54
$A4 = A6 * A9 + A8 * A9$	$A4 = 681.94 * A6 * A9 + 127.34 * A8 * A9$	81.58

Table 7: The conversion functions generated for Set 2

$A5$	Conversion function	Support (%)
0	$A4 = 400 * A6 * A9$	100
1	$A4 = 700 * A6 * A9$	100
2	$A4 = 1000 * A6 * A9$	100
3	$A4 = 1800 * A6 * A9$	100
4	$A4 = 5800 * A6 * A9$	100

Table 8: The conversion functions generated for Set 3

$A5$	Conversion function	Support (%)
0	$A2 = 0.4 * A9$	100
1	$A2 = 0.7 * A9$	100
2	$A2 = 1.0 * A9$	100
3	$A2 = 1.8 * A9$	100
4	$A2 = 5.8 * A9$	100

Table 9: Attributes of sales data sample

Attribute	Data Source	Name	Description
A1	T(ransaction)	month	financial period
A2	T(ransaction)	invoice_no	the unique invoice number
A3	T(ransaction)	amount	total amount invoiced in the original currency
A4	T(ransaction)	sales_type	type of sales that determines tax
A5	T(ransaction)	GST_rate	goods and service tax rate
A6	T(ransaction)	currency	currency code
A7	T(ransaction)	exchange_rate	exchange rate for the invoice
A8	T(ransaction)	GST_amount	goods and service tax in the original currency
A9	C(ost/profit)	amount	amount for the sales account, i.e. before-tax amount
A10	A(ccounts)	amount	amount of the invoice in the local currency

Table 10: Partial correlation analysis between A10 and {A3,A5,A7,A8}

	Correlation Efficient of A10 with		
	Zero-order	Controlling A3	Controlling A8
A3	0.91778708	-	0.8727613
A5	0.02049382	0.1453855	-0.2968708
A7	0.02492715	0.4562510	-0.006394451
A8	0.71552943	0.5122901	-

Table 11: Partial correlation analysis between A9 and {A3,A5,A7,A8}

	Correlation Efficient of A9 with		
	Zero-order	Controlling A3	Controlling A8
A3	0.91953641	-	0.8739174
A5	0.01350388	0.1292704	-0.2976056
A7	0.02355110	0.4581783	-0.007597317
A8	0.70452050	0.4791232	-

Table 12: Results after partitioning data using *sales.type*

$A4$	No	Conversion functions	Support	Number of Outliers
1	1	$A9 = 0.9708738 * A3 * A7$	99.94	2
	2	$A9 = 33.33333 * A7 * A8$	99.36	21
2	3	$A9 = A3 * A7$	100.00	0

Table 13: Two erroneous tuple discovered

month	inv-no	amt.	GST type	GST rate	currency	exch. rate	GST	C.amount	A.amount
3	8237	45.50	1	3	1	1.00	0.00	45.50	45.50
6	12991	311.03	1	3	1	1.00	6.53	304.50	311.03

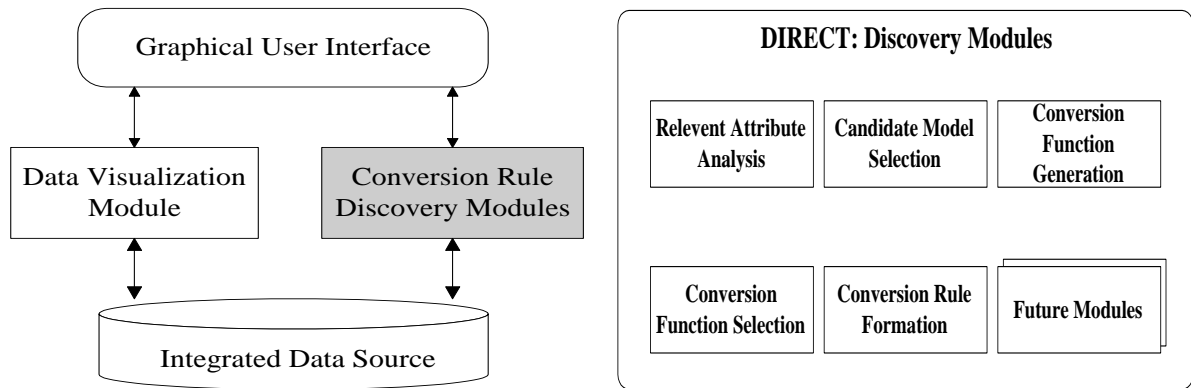


Figure 1: **DIRECT**: The reference architecture and conversion rule discovery modules

Rule Mining




Figure 2: Discovering data value conversion rules from data




Figure 3: A snapshot of the output of DIRECT visualisation module




Figure 4: Output of relevant attribute analysis and their scatter plot

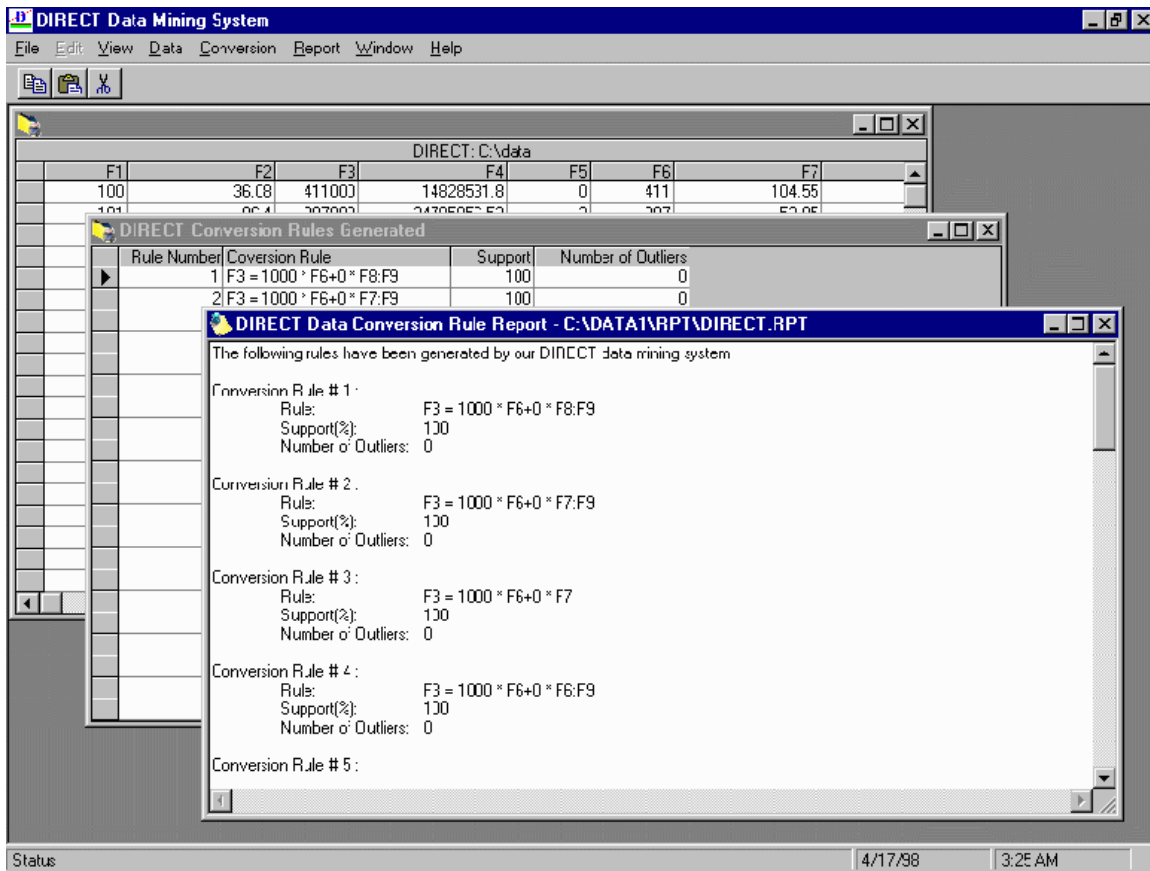


Figure 5: Conversion rule generation result