



How to Semantically Enhance a Data Mining Process?

Laurent Brisson, Martine Collard

► **To cite this version:**

Laurent Brisson, Martine Collard. How to Semantically Enhance a Data Mining Process?. Lecture Notes in Business Information Processing, Springer-Verlag, 2009, 13, pp.103-116. hal-00476279

HAL Id: hal-00476279

<https://hal.archives-ouvertes.fr/hal-00476279>

Submitted on 26 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How to Semantically Enhance a Data Mining Process?

Laurent Brisson^{1,3} and Martine Collard^{2,4}

¹ Institut TELECOM, TELECOM Bretagne, CNRS UMR 3192 LAB-STICC

Technopôle Brest-Iroise CS 83818, 29238 Brest Cedex 3, France

laurent.brisson@telecom-bretagne.eu

<http://perso.enst-bretagne.fr/laurentbrisson/>

² INRIA Sophia Antipolis, 2004 route des Lucioles, 06902 BP93 Sophia Antipolis, France

martine.collard@sophia.inria.fr

³ Université européenne de Bretagne, France

⁴ Université Nice Sophia Antipolis, France

Abstract. This paper presents the KEOPS data mining methodology centered on domain knowledge integration. KEOPS is a CRISP-DM compliant methodology which integrates a knowledge base and an ontology. In this paper, we focus first on the pre-processing steps of business understanding and data understanding in order to build an ontology driven information system (ODIS). Then we show how the knowledge base is used for the post-processing step of model interpretation. We detail the role of the ontology and we define a part-way interestingness measure that integrates both objective and subjective criteria in order to eval model relevance according to expert knowledge. We present experiments conducted on real data and their results.

Keywords: Data mining, Knowledge integration, Ontology Driven Information System.

1 Introduction

In knowledge discovery from data, methods and techniques are developed for discovering specific trends in a system or organization business by analyzing its data. The real advantage for decision making relies on the add-on provided by comparing extracted knowledge against the *a priori* knowledge of domain expert. Integrating domain *a priori* knowledge during the data mining process is currently an important research issue in the data mining field.

In this paper, we present the KEOPS methodology based on an ontology driven information system which integrates *a priori* knowledge all along the data mining process in a coherent and uniform manner. We detail each of these ontology driven steps and we define a part-way interestingness measure that integrates both objective and subjective criteria in order to evaluate model relevance according to expert knowledge.

The paper is organized in six sections. Section 2 presents related works. Section 3 presents the KEOPS methodology step by step. In Section 4, we comment some results which demonstrate the relevance of the approach. We discuss our approach in section 5 and conclude in section 6.

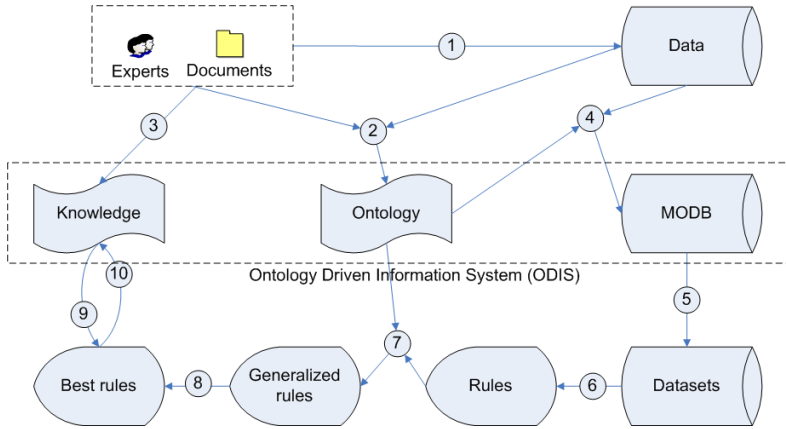


Fig. 1. KEOPS methodology

2 Related Works

2.1 Knowledge Integration in Data Mining

The Data Mining process described according to the CRISP-DM model [1] is presented as both iterative and interactive. The iterative nature is due to the way processes run cycling test-error experiments. Indeed data miners have to repeat the pre-processing steps of domain understanding, data understanding and data preparation until final models are considered relevant. The interactive nature is inherent to a data mining activity since communications with experts is necessary for understanding domain and data and for interpreting results. Issues in evaluating and interpreting mining process results are currently big research challenges. In order to avoid useless iterations on preliminary tasks and facilitate model interpretation, one solution is to explore deeply expert knowledge and source data in order to formalize them in conceptual structures and exploit these structures both for robust data preparation and for flexible model interpretation.

In the literature, partial solutions for domain knowledge interpretation are proposed for optimizing pre-processing steps [2]. For model evaluation, detailed studies have been devoted to interestingness measures [3]. A consensus among researchers is now established to consider objective interestingness versus subjective interestingness. Objective interestingness is traditionally evaluated by a variety of statistic indexes while subjective interestingness is generally evaluated by comparing discovered patterns to user knowledge or *a priori* convictions of domain experts. In this paper we present the KEOPS methodology based on an ontology driven information system which addresses the knowledge integration issue (see figure 1). The system relies on three main components: an ontology, a knowledge base and a mining oriented database rebuilt from source raw data. These components allow to model domain concepts and relationships among them. They are used to pre-process data and to identify mapping between discovered patterns and expert knowledge.

2.2 Ontology Driven Information System (ODIS)

An ontology driven information systems is an information system (IS) which relies mainly on an *explicit* ontology. This ontology may underlie all aspects and components of the information system. An ODIS contains three kinds of components: application programs, information resources and user interfaces. [4] discusses the impact of an ontology on an information system according to temporal and structural dimension.

The temporal dimension refers to ontology role during IS construction and run-time. If we have a set of reusable ontologies, the semantic content expressed can be transformed and translated into an IS component. Even if the volume of ontology knowledge available is modest it may nevertheless help a designer in a conceptual analysis task. This task consists frequently of redesigning an existing information system. This approach fits the needs of data mining tasks where an operational database has to be transformed into datasets before the data mining modeling step.

The structural dimension refers to each information system component which may use the ontology in a specific way.

- Database component: at development time, an ontology can play an important role in requirement analysis and conceptual modeling. The resulting conceptual model can be represented as a computer processable ontology mapped to a concrete target platform [5]. Usually, IS conceptual schemes (CS) are created from scratch, wasting a lot of time and resources.
- Interface components may be assisted by ontologies which are used to generate personalized interfaces or to manage user profiles [6,7].
- Application program components use implicit knowledge in order to perform a task. However, this knowledge is often hardcoded in software. Ontologies may provide a formal base helping to access domain knowledge.

2.3 Ontology-Based Validation Methods

Subjective interestingness measure were developed in order to complement objective measure and give an insight on real human interest. However, these measures lack of semantic formalization, and force the user to express all of his expectations. Consequently, extracted pattern validation process must involve not only the study of patterns but also the use of a domain ontology and domain experts expectations. Rules expressed to filter out noisy pattern or to select the most interesting ones should be relevant.

An important issue in ontology-based validation methods is the definition of semantic similarity measures between ontologies concepts. Fortunately, there are numerous works that address this problem. We can consider two kinds of methods in order to measure semantic similarity within an ontology: *edge counting* methods and *information-theoretic* based methods.

Edge counting methods consist of calculating the distance between ontology concepts, similarity decreasing while distance increasing. If there are several paths, minimum or average distances can be used. Leacock and Chodorow [8] measure semantic similarity by finding the shortest method distance between two concepts and then scale the distance by the maximum distance in the "is-a" hierarchy. Choi and Kim[9], use hierarchy concepts tree to calculate a concept distance between two concepts. Zhong and

al. [10] define weights for the links among concepts according to their position in the taxonomy.

Resnik introduced *information-theoretic* measures [11,12] based on the information content of the lower common ancestor of two concepts. The information content of a term decreases with its occurring probability. If the lower common ancestor of two concepts is a generic concept, these concepts should be pretty different and their lower common ancestors have a low information level. Resnik demonstrated that such *information-theoretic* based methods are less sensitive, and in some cases not sensitive, to the problem of link density variability [11]. Lin [13] improves Resnik's measure considering how close are the concepts to their lower common ancestor. Jiang presents a combined approach that inherits the *edge counting* based approach and enhanced it by node-based approach of the information content calculation [14]. Lord compared Resnik's, Lin's and Jiang's measures in order to use them to explore the Gene Ontology (GO). His results suggest that all three measures show a strong correlation between sequence similarity and molecular function semantic similarity. He concludes that none of the three measures has a clear advantage over the others, although each one has strengths and weaknesses [15]. Schlicker and al. [16] introduced a new measure of similarity between GO terms that is based on Resnik's and Lin's definitions. This measure takes into account how close these terms are to their lower common ancestor and uses a *score* allowing to identify functionally related gene products from different species that have no significant sequence similarity.

In KEOPS, we introduce ontology-based post-processing and evaluation steps too.

3 KEOPS Methodology

KEOPS is a methodology which drives data mining processes by integrating expert knowledge. These are the goals addressed:

- To manage interactions between knowledge and data all along the data mining process: data preparation, datasets generation, modeling, evaluation and results visualization.
- To evaluate extracted models according to domain expert knowledge.
- To provide easy navigation throughout the space of results.

KEOPS (cf. fig. 1) is based upon an ontology driven information system (ODIS) set up with four components:

- An application ontology whose concepts and relationships between them are dedicated to domain and data mining task.
- A Mining Oriented DataBase (MODB): a relational database whose attributes and values are chosen among ontology concepts.
- A knowledge base to express consensual knowledge, obvious knowledge and user assumptions.
- A set of information system components - user interfaces, extraction algorithms, evaluation methods - in order to select the most relevant extracted models according to expert knowledge.

KEOPS methodology extends the CRISP-DM process model by integrating knowledge in most steps of the mining process. The initial step focuses on business understanding. The second step focuses on data understanding and activities in order to check data reliability. Data reliability problems are solved during the third step of data preparation. The fourth step is the evaluation of extracted models. In this paper we don't focus on modeling step of CRISP-DM model since we ran CLOSE algorithm [17] which extracts association rules without domain knowledge.

3.1 Business Understanding

During business understanding step, documents, data, domain knowledge and discussion between experts lead to assess situation, to determine business objectives and success criteria, and to evaluate risks and contingencies. However this step is often rather informal.

KEOPS methodology requires to build an ontology driven information system during the next step, data understanding. Consequently an informal specification of business objectives and expert knowledge is henceforth insufficient. Thus, it is necessary to formalize expert knowledge during business understanding. We chose to state knowledge with production rules, also called "if ... then ..." rules. These rules are modular, each defining a small and independent piece of knowledge. Furthermore, they can be easily compared to extracted association rules. Each knowledge rule has some essential properties to select the most interesting association rules:

- Knowledge confidence level: five different values are available to describe knowledge confidence according to a domain expert. These values are ranges of confidence values: 0-20%, 20-40%, 40-60%, 60-80% and 80-100%. We call confidence the probability for the rule consequence to occur when the rule condition holds.
- Knowledge certainty:
 - Obvious: knowledge cannot be contradicted.
 - Consensual: domain knowledge shared among experts.
 - Assumption: knowledge the user wants to check.

Since the description of expert interview methodology in order to capture knowledge is beyond the scope of this paper, the reader should refer to [18].

3.2 Data Understanding

Data understanding means selection and description of source data in order to capture their semantic and reliability. During this step, the ontology is built in order to identify domain concepts and relationships between them (the objective is to select among data the most interesting attributes according to the business objectives), to solve ambiguities within data and to choose data discretization levels.

Consequently, the ontology formalizes domain concepts and information about data. This ontology is an application ontology; it contains the essential knowledge in order to drive data mining tasks. Ontology concepts are related to domain concepts, however relationships between them model database relationships. During next step, data preparation (cf. section 3.3), a relational database called Mining Oriented DataBase (MODB) will be built.

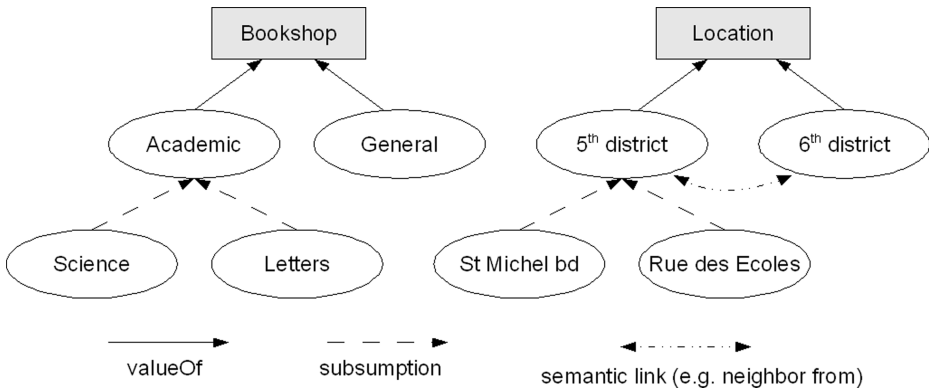


Fig. 2. Bookshop ontology snapshot

In order to understand links between the MODB and the ontology it is necessary to define notions of domain, concept and relationships:

- Domain: This notion in KEOPS methodology, refers to the notion of domain in relational theory. A domain represents a set of values associated to a semantic entity (or concept).
- Concept: Each concept of the ontology has a property defining its role. There exist two classes of concepts: attribute concepts and value concepts.
 - An *attribute concept* is identified by a name and a *domain*.
 - Each value of domain is called a *value concept*.

Thus a domain is described by an attribute concept and by value concepts organized into a taxonomy. Each MODB attribute is linked to one and only one attribute concept and takes its values in the associated domain. In figure 2 “Bookshop” is an attribute concept, “Academic” a value concept and the set {Academic, General, Sciences, Letters} defines “Bookshop” domain.

- Relationships: There exists three kinds of relationships between concepts:
 - A data-related relationship: “valueOf” relationship between an attribute concept and a value concept. The set of value concepts linked to an attribute concept with “valueOf” relationship define a domain within the MODB.
 - Subsumption relationship between two value concepts. A concept subsumed by another one is member of the same domain. This relationship is useful during data preparation (to select data granularity in datasets), reduction of rule volume (to generate generalized association rules, see 3.4, comparison between models and knowledge (to consider sibling and ancestor concepts) and final results visualization.
 - Semantic relationships between value concepts. These relationships could be order, composition, exclusion or equivalence relationships. They can be used to compare extracted models and knowledge and to visualize results.

KEOPS methodology aims to extract interesting models according user knowledge. Consequently, it is necessary during ontology construction to be careful with some usual concerns in data mining:

- Aggregation level: like data, ontology concepts have to represent disjoint domains.
- Discretization level: ontology concepts have to model various solutions for data discretization. Bad choices may affect modeling step efficiency.
- Data correlation: if concepts are strongly related into the MODB, extracted models might be trivial and uninteresting.

Since these concerns are beyond the scope of this paper, the reader should refer to [19] for a better insight on concept elicitation and [20] for a better insight on discretization and grouping.

Table 1. Bookshop ontology concept elicitation

| Source Data | Attribute Concept | Value Concept |
|-----------------|-------------------|---------------|
| St Denis Shop | Bookshop | Academic |
| St Denis Shop | Location | St Michel bd |
| Rive Gauche 5th | Bookshop | General |
| Rive Gauche 5th | Location | 5th District |

Example. Let's take the case of a bookstore company with several bookshops in Paris which plan to improve customer relationships. Bookshops may be specialized in a field like "academic" or not (general) (see figure 2). Bookshops are located geographically. Data are provided on bookshops, customers and sales. Table 1 shows a way for mapping source values to ontology concepts.

3.3 Data Preparation

Data preparation is very iterative and time consuming. The objective is to refine data: discretize, clean and build new attributes and values in the MODB. During this step, KEOPS suggests building MODB by mapping original data with ontology concepts. The database contains only bottom ontology concepts. The objective is to structure knowledge and data in order to process efficient mining tasks and to save time spent into data preparation. The idea is to allow generation of multiple datasets from the MODB, using ontology relationships without another preparation step from raw data. Furthermore, during ODIS construction, experts can express their knowledge using the ontology which is consistent with data.

Mining Oriented Database (MODB) Construction. Databases often contain several tables sharing similar information. However, it is desirable that each MODB table contains all the information semantically close and it's important to observe normal forms in these tables. During datasets generation, it's easy to use join in order to create interesting datasets to be mined. However these datasets don't have to observe normal forms.

Datasets Generation. It’s often necessary, in a data mining process, to step back to data preparation. Algorithms were proposed to choose relevant attributes among large data sources. However, sometimes results don’t satisfy user expectations and datasets have to be built again to run new tests. KEOPS methodology suggests using the ontology in order to describe domain values and relationships between these values. Consequently, various datasets could be generated according to expert user choices.

The ontology driven information system allows choosing all data preparation strategies providing various datasets from the same source values. A dataset is built using the following operators:

- Traditional relational algebra operators: projection, selection and join.
- Data granularity: this operator allows choosing, among ontology, concepts which will be in the mining oriented database.

In order to generate datasets we developed software whose inputs are MODB and user parameters and outputs are new datasets. The user can graphically select relational algebra operator and data granularity. While database attributes and values are also ontology concepts KEOPS methodology and KEOPS software make easier the data preparation task.

3.4 Evaluation

This step assesses to what extent models meet the business objectives and seeks to determine if there is some business reason why these models are deficient. Furthermore, algorithms may generate lots of models according to parameters chosen for the extraction. That’s why *evaluation* is an important task in KEOPS methodology in order to select the most interesting models according to expert knowledge.

Table 2. Interestingness measure if confidence levels are similar

| Kind of knowledge | Rule R informative level | | |
|-------------------|--------------------------|---------|-------------|
| | More than K | Similar | Less than K |
| Obvious | weak | none | none |
| Consensual | medium | weak | weak |
| Assumption | strong | medium | medium |

Rule Volume Reduction. We choose an association rule extraction algorithm which can generate bases containing only minimal non-redundant rules without information loss. Then, these rules are filtered to suppress semantic redundancies. KEOPS methodology is based on Srikant’s *generalized association rules* definition [21]. These rules are minimal because they forbid all irrelevant relationships within their items. We give a formal definition below:

Let \mathcal{T} be a taxonomy of items. $R : A \rightarrow C$ is called *generalized association rule* if:

- $A \subset \mathcal{T}$
- $C \subset \mathcal{T}$
- No item in C is an ancestor of any item in A or C
- No item in A is an ancestor of any item in A

Consequently relationships appearing within these rules are semantic and generalization relationships from C items to A items. The objective is to maximize information level in minimal rules. The last step consists of replacing a set of these rules by a more generalized one.

Rule Interestingness Evaluation. KEOPS methodology suggests comparing extracted rules with expert's knowledge. Extracted rules having one or more items that are in relationship with some knowledge rules items (i.e. value concepts are linked in the ontology) have to be selected. Then, for each pair knowledge rule/extracted rule:

- Extracted rule antecedant coverage is compared to knowledge rule antecedent coverage, then extracted rule consequent coverage is compared to knowledge rule consequent coverage.
- By coverage comparison the most informative rule is deduced, i.e. rule predicting the largest consequence from the smallest condition.
- IMAK interestingness measure is applied [22]. This measure evaluates rule quality considering relative confidence values, relative information levels and knowledge certainty (see section 3.1).

Thus, ontology driven information systems are useful in order to formalize domain concepts, to express knowledge, to generate models and to facilitate knowledge and models ontology-based comparison.

Example. Let us assume that a domain expert makes the following assumption: “*If a student wants to buy a book about JAVA he comes to an academic bookshop.*” and gives it a 60%-80% estimation of confidence. Let us assume that the extracted rule is slightly different because it says that “*Every young customer buying a book about J2EE comes to an academic bookshop*” and has 75% confidence.

Assumption K $\text{book}=\text{'JAVA'} \wedge \text{buyer}=\text{'student'} \rightarrow \text{bookshop}=\text{'Academic'}$

Extracted Rule R $\text{book}=\text{'J2EE'} \wedge \text{buyer}=\text{'youngs'} \rightarrow \text{bookshop}=\text{'Academic'}$

According KEOPS methodology these two rules are said to be comparable because at least one extracted rule item is in relationship with a knowledge rule item: 'youngs' is more general than 'student' and 'JAVA' is more general than 'J2EE'. Then, the algorithm compares the coverage of these two rules in order to evaluate the more informative one. Let's make the assumption that R is more informative than K . Since these two rules have similar confidence we can use table 2 in order to evaluate extracted rule interestingness (similar tables for various confidence levels are presented in [22]). While the knowledge is an assumption, the interestingness degree of the extracted rule is **strong**.

4 Experiments

Although we illustrated in this paper the KEOPS methodology with bookstore example, we run experiments on real data provided by French Family Allowance Office (CAF:


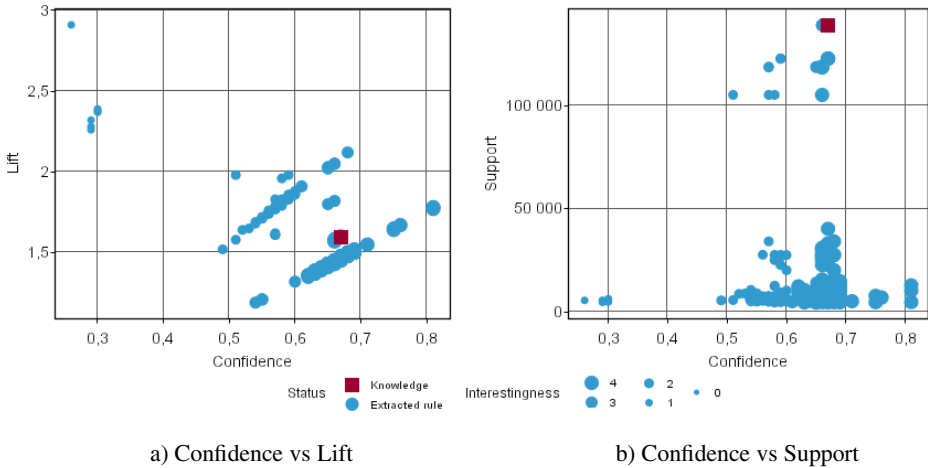


Fig. 3. Confidence vs Lift of all of the extracted rules



a) Confidence vs Lift

b) Confidence vs Support

Fig. 4. Extracted rules (dots) matching knowledge rule 335 (square) (IMAK interestingness value increase with dot size)

Caisses d’allocations familiales). In this section we don’t express some specific rules about allowance beneficiaries behavior (because of privacy) but only extracted rules reliability. These results show we are able to select relevant rules to provide to experts for final human evaluation.

CAF data were extracted during 2004 in the town of Grenoble (France). Each row describes one contact between the office and a beneficiary with 15 attributes and data about 443716 contacts were provided. We ran CLOSE algorithm and extracted 4404 association rules. The interestingness measure, IMAK, helps to filter the best ones. Figure 3 plots 4404 rules according to confidence and lift.

Experiments illustrated by figure 4 and 5 compare these rules to a specific knowledge. We may observe that among all of the extracted rules only few of them are




Fig. 5. Extracted rules (dots) matching knowledge rule 565 (square) (IMAK interestingness value increase with dot size)

selected. Selection condition is to match the knowledge and to have an interestingness value greater than 0. In these figures interestingness value is illustrated by the dot size.

In figure 4 lift of selected rules is greater than 1 and often greater than knowledge lift (lift equals 1 at independency). Furthermore, some extracted rules have a better confidence but a smaller support: they illustrated the discovery of rare events which could be very interesting for expert users.

Figure 5 shows some results for another specific knowledge. We may observe again that only few rules are selected. These rules offer various tradeoff to select rare events (low support and high confidence) or general rules (high support and good confidence) to provide to domain experts.

5 Discussion

As future work, we plan to evaluate rules selected by KEOPS system on a larger scale with the help of some expert groups who are able to validate their semantic relevance. However we stand up to the problem of pattern management in a coherent way with knowledge management. These patterns are heterogeneous (decision trees, clusters, rules etc.) and it is a laborious task to access and analyze them. Researches on pattern management aims at setting up systems to maintain the persistence and availability of results. They have to represent patterns for sharing and reasoning and to manage them in order to allow efficient searches among them.

Consequently there is a need for services to store patterns (indexation of various patterns), to manage patterns (insert, delete, update patterns), to manage metadata (link to data sources, temporal information, semantic information, quality measures) and to query patterns (to retrieve efficiently patterns with some constraints, to evaluate similarity).

Several approaches defined logical models for pattern representation (PMML, CWM-DM, etc.). Although they are well-suited for data models sharing they seem inadequate

to represent and manage patterns in a flexible and coherent way. Inductive databases (IDB) [23] provide models for pattern storage and inductive query languages. These languages allow to generate patterns satisfying some user-specified constraints (using data-mining algorithm) and to querying previously generated patterns. Rizzi and al. [24] defined Pattern Base Management System (PBMS) claiming that a logical separation between database and pattern-base is needed to ensure efficient handling of both raw data and patterns through dedicated management systems.

In order to extend our works by enhancing post-processing steps with expert knowledge we plan to develop a system dedicated to pattern storage and querying. In a first stage, we'll focus only on rules based patterns: association rules, classification rules or sequential rules. The KEOPS current rule interestingness evaluation method is mainly based on the previously defined IMAK measure [22]. We plan to enrich the method according to different ways:

- Rather than confronting one extracted rule according each of the existing knowledge rules in the base, we will measure the rule relative interestingness by comparison with a set of corresponding knowledge rules. Assuming that the knowledge rule set is relevant and does not contain any contradiction, the interestingness of a given rule will be evaluated relatively to a set of similar knowledge rules since it may either generalize them or highlight a more precise case than them.
- The current measure is computed partly upon rule quality indices as support, confidence and lift. These indices are easily interpretable. They allow to select rules according to the common sense of coverage and implication but they are quite simple. More sophisticated interestingness measures could be used and combined to emphasize precise objectives of the user [25] and it will be valuable to observe how extracted rules are different.
- The ODIS is the backbone of the KEOPS approach and one of our further works will more deeply take advantage of the semantic relationships stored among ontological concepts. Semantic distances [26,11,10] as discussed in section 2.3 will be integrated in order to determine more accurately similarity measures between an extracted rule and a knowledge rule and to refine the interestingness measure too.

6 Conclusions

Managing domain knowledge during the data mining process is currently an important research issue in the data mining field. In this paper, we have presented the so-called KEOPS methodology for integrating expert knowledge all along the data mining process in a coherent and uniform manner.

We built an ontology driven information system (ODIS) based on an application ontology, a knowledge base and a mining oriented database re-engineered from source raw data. Thus, expert knowledge is used during business and data understanding, data preparation and model evaluation steps. We have shown that integrating expert knowledge during the first step, gives experts a best insight upon the whole data mining process. In the last step we have introduced IMAK, a part-way interestingness measure that integrates both objective and subjective criteria in order to evaluate models relevance according to expert knowledge.

We implemented KEOPS prototype in order to run experiments. Experimental results show that IMAK measure helps to select a reduced rule set among data mining results. These rules offer various tradeoff allowing experts to select rare events or more general rules which are relevant according to their knowledge.

References

1. Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., Wirth, R.: *Crisp-dm 1.0: Step-by-step data mining guide*. In: SPSS Inc. (2000)
2. Kedad, Z., Métais, E.: *Ontology-based data cleaning*. In: Andersson, B., Bergholtz, M., Johannesson, P. (eds.) *NLDB 2002*. LNCS, vol. 2553, pp. 137–149. Springer, Heidelberg (2002)
3. McGarry, K.: *A survey of interestingness measures for knowledge discovery*. *Knowl. Eng. Rev.* 20, 39–61 (2005)
4. Guarino, N.: *Formal Ontology in Information Systems*. IOS Press, Amsterdam (1998); Amended version of previous one in *Proceedings of the 1st International Conference, Trento, Italy, June 6-8 (1998)*
5. Ceri, S., Fraternali, P.: *Designing Database Applications with Objects and Rules: The IDEA Methodology*. Series on Database Systems and Applications. Addison-Wesley, Reading (1997)
6. Guarino, N., Masolo, C., Vetere, G.: *Ontoseek: Using large linguistic ontologies for gathering information resources from the web*. Technical report, LADSEB-CNR (1998)
7. Penarrubia, A., Fernandez-Caballero, A., Gonzalez, P., Botella, F., Grau, A., Martinez, O.: *Ontology-based interface adaptivity in web-based learning systems*. In: *ICALT 2004: Proceedings of the IEEE International Conference on Advanced Learning Technologies (ICALT 2004)*, Washington, DC, USA, pp. 435–439. IEEE Computer Society, Los Alamitos (2004)
8. Leacock, C., Chodorow, M.: *Combining local context with wordnet similarity for word sense identification*. In: Fellbaum, C. (ed.) *WordNet: A Lexical Reference System and its Application*. MIT Press, Cambridge (1998)
9. Choi, I., Kim, M.: *Topic distillation using hierarchy concept tree*. In: *ACM SIGIR conference*, pp. 371–372 (2003)
10. Zhong, J., Zhu, H., Li, J., Yu, Y.: *Conceptual graph matching for semantic search*. In: *ICCS conference*, pp. 92–196 (2002)
11. Resnik, P.: *Using information content to evaluate semantic similarity in a taxonomy*. In: *IJCAI conference*, pp. 448–453 (1995)
12. Resnik, P.: *Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language*. *Journal of Artificial Intelligence Research* 11, 95–130 (1999)
13. Lin, D.: *An information-theoretic definition of similarity*. In: *ICML conference* (1998)
14. Jiang, J., Conrath, D.W.: *Semantic similarity based on corpus statistics and lexical taxonomy*. *CoRR cmp-lg/9709008* (1997)
15. Lord, P., Stevens, R., Brass, A., Goble, C.A.: *Semantic similarity measures as tools for exploring the gene ontology*. In: *PSB conference* (2003)
16. Schlicker, A., Domingues, F., Rahnenfuhrer, J., Lengauer, T.: *A new measure for functional similarity of gene products based on gene ontology*. *BMC Bioinformatics* 7, 302 (2006)
17. Pasquier, N., Bastide, Y., Taouil, R., Lakhil, L.: *Closed set based discovery of small covers for association rules*. In: *Actes des 15èmes journées Bases de Données Avancées (BDA 1999)*, pp. 361–381 (1999)
18. Becker, H.S.: *Sociological Work: Method and Substance*. Transaction Publishers, U. S (1976)

19. De Leenheer, P., de Moor, A.: Context-driven disambiguation in ontology elicitation. In: Shvaiko, P., Euzenat, J. (eds.) *Context and Ontologies: Theory, Practice and Applications*, Pittsburgh, Pennsylvania, AAAI, pp. 17–24. AAAI Press, Menlo Park (2005)
20. Berka, P., Bruha, I.: Discretization and grouping: Preprocessing steps for data mining. In: Żytkow, J.M. (ed.) *PKDD 1998*. LNCS, vol. 1510, pp. 239–245. Springer, Heidelberg (1998)
21. Srikant, R., Agrawal, R.: Mining generalized association rules. In: *VLDB 1995: Proceedings of the 21th International Conference on Very Large Data Bases*, pp. 407–419. Morgan Kaufmann Publishers Inc., San Francisco (1995)
22. Brisson, L.: Knowledge extraction using a conceptual information system (ExCIS). In: Collard, M. (ed.) *ODBIS 2005/2006*. LNCS, vol. 4623, pp. 119–134. Springer, Heidelberg (2007)
23. Imieliński, T., Mannila, H.: A database perspective on knowledge discovery. *Commun. ACM* 39, 58–64 (1996)
24. Rizzi, S., Bertino, E., Catania, B., Golfarelli, M., Halkidi, M., Terrovitis, M., Vassiliadis, P., Vazirgiannis, M., Vrachnos, E.: Towards a logical model for patterns. In: Song, I.-Y., Liddle, S.W., Ling, T.-W., Scheuermann, P. (eds.) *ER 2003*. LNCS, vol. 2813, pp. 77–90. Springer, Heidelberg (2003)
25. Collard, M., Vansnick, J.C.: How to measure interestingness in data mining: a multiple criteria decision analysis approach. In: *RCIS*, pp. 395–400 (2007)
26. Rada, R., Mili, H., Bicknell, E., Blettner, M.: *Development and application of a metric on semantic nets*, vol. 19, pp. 17–30 (1989)