Towards Knowledge Discovery in the Semantic Web

Thomas Fischer, Johannes Ruhland

Department of Information Systems,
Friedrich Schiller University Jena

1 Introduction

In the past, data mining and machine learning research has developed various techniques to learn on data and to extract patterns from data to support decision makers in various tasks, such as customer profiling, targeted marketing, store layout, and fraud detection (Tan et al., 2005, p.1). In addition, the World Wide Web increasingly offers distributed information that can be useful for strategic, tactical or operational decisions, including news, events, financial information, information about competitors as well as information about the social networks of customers and employees etc. The Web thus has the potential for a high impact on competitive actions and competitive dynamics of enterprises that should utilize this information. However, the growing amount of these distributed information resources leads to a dilemma:"... *the more distributed and independently managed that resources on the Web become, the greater is their potential value, but the harder it is to extract value…*" (Singh and Huhns, 2005, p.7). On the one hand the human ability for information processing is limited (Edelmann, 2000, p.168), whilst otherwise the amount of available information of the Web increases exponentially, which leads to increasing information saturation (Krcmar, 2004, p.52). In this context, it becomes more and more important to detect useful patterns in the Web, thus use it as a rich source for data mining (Berendt et al., 2002; Han and Kamber, 2006, p.628) in addition to company internal databases.

The extraction of information and interesting patterns out of the Web is a complex task, because the current Web is mainly utilized for human consumption. This means that the available information is represented by mark-up languages such as XHTML[1] and its predecessors that describe only a visual presentation. Unfortunately, these languages are not sufficient to let software agents "under-

[1] http://www.w3.org/TR/xhtml1/

stand" the information they are processing. For instance, the character string "Jena" does neither reflect to a machine that this is the name of a city[2], nor does it reflect that this is also the title of a famous semantic web framework[3]. Due to this ambiguity, the discovery of useful patterns in such unstructured information is very difficult and has been addressed by research on web mining (Stumme et al., 2006).

However, there have been also increasing efforts in the research community to realize the vision of the so called Semantic Web: "*The Semantic Web is not a separate Web but an extension of the current one, in which information is given well-defined meaning, better enabling computers and people to work in cooperation*" (Berners-Lee et al., 2001). It seems therefore to be valuable to perform data mining on information with a well-defined meaning to improve the knowledge discovery process.

The utilization of data mining on semantic web information for business intelligence has got not much attendance in the research community in comparison to the overall research investments in this field. Furthermore, there are a lot of open topics that have to be addressed. In this paper we motivate this field of research by a scenario to outline the differences of the knowledge discovery process as well as to deduce requirements.

The remainder of this paper is organized as follows. Section 2 outlines the research context. Section 3 describes a scenario for relational association rules, which serves as a basis for an overview about the requirements for knowledge discovery in the semantic web. Finally, Section 4 concludes this paper.

2 Qualifying the Research Context

2.1 Semantic Web

The Semantic Web (Berners-Lee et al., 2001) focuses on the extension of the current Web by machine readable and "understandable" meta information. The vocabulary of these statements is typically derived from one or more ontologies, which are a shared conceptualization of the domain of discourse (Gruber, 1993). The semantic description (meaningful to a machine) of Web data has been driven by the research community through the creation of different standards, for instance, the Resource Description Framework (RDF) (Klyne et al., 2004), the Resource Description Framework Schema RDF(S) (Brickley et al., 2004) and the Web Ontology Language (OWL) (Smith et al., 2004).[4] These approaches provide a formal way to specify shared vocabularies that can be used in statements about resources. Furthermore, they utilize a syntax based on the Extensible Markup Lan-

[2] http://www.jena.de/

[3] http://jena.sourceforge.net/

[4] There are also other ontology languages, but the mentioned standards are widely accepted in the research literature. The proposed approach is independent from the ontology language as long as the language is based on description or first-order-logics.

guage (XML) (Beckett and McBride, 2004) and thus can be effectively processed by machines.

Resource Description Framework (RDF) and RDF-Schema

The Resource Description Framework (RDF) (Klyne et al., 2004) is a framework for representing information on the Web. RDF allows anyone to make statements about any resource, which could be a material or immaterial thing. A statement is defined as a triple, consisting a subject s, predicate p and object o, written as $p(s,o)$. This means that s has a predicate (or property) p with value o. RDF is based on a graph data model. A RDF graph $G=(V,E)$ is a representation of the document triples. A node n of the graph could be a subject s or object o, which is connected through a directed arc (s,o) that represents the predicate p.

RDF-Schema (Brickley et al., 2004) is a minimal ontological language. It has capabilities to define classes and properties, and enables the specification of how they should be used together. Classes and properties could be arranged in a hierarchy. Instances of a class are referenced to its class through the "*rdf:type*" definition. RDF-Schema provides means to define a simple shared vocabulary. Nevertheless, its expressiveness is limited. Amongst others things, it provides no support for cardinality constraints on properties, transitive properties as well as equivalence and disjointness relationships of classes and individuals. The Web Ontology Language (OWL) is more expressive than RDF-Schema and is thus considered below in more detail.

Web Ontology Language (OWL)

The Web Ontology Language (OWL) (Smith et al., 2004) is build on top of RDF and RDF-Schema. OWL provides the three sub languages OWL-Lite, OWL-DL and OWL-Full. The usage of a language depends on the needed expressiveness of the ontology (Maedche et al., 2003). OWL-Lite and OWL-DL are widespread used sub languages that are based on the formalisms of description logics (DL).

The logical structure of a DL knowledge base is based on a so called *TBox* and *ABox*: $KB = (TBox, ABox)$. The *TBox* contains intensional knowledge representation and is build through the definition of concepts and properties. The *ABox* contains assertions about the named individuals in terms of the defined vocabulary. Furthermore, the *ABox* depends on the current circumstances and is part of constant change. A detailed overview about description logics can be found in Baader et al. (2003). In this paper we motivate the utilization of relational data mining techniques to discover patterns in information with such a well-defined meaning.

2.2 Data Mining

Data mining (Tan et al., 2005) refers to extracting valid, novel, potentially useful, and ultimately understandable patterns from large amounts of data and is part of the KDD (Knowledge Discovery from Data) process (Han and Kamber, 2006; Fayyad et al., 1996). The process of knowledge discovery is based on *data cleaning*, *data integration*, *data selection*, *data transformation*, *data mining*, *pattern evaluation* and *knowledge representation*.

In the data mining step different techniques such as cluster analysis, predictive modelling (classification - discret / regression - continuous), association analysis, anomaly detection, summarization, evolution analysis can be utilized to achieve the aims of a specific knowledge discovery process.

As stated above, web resources can hold a lot of useful information and it is therefore interesting to apply data mining techniques on them, which is called web mining. "*Web mining is the application of data mining techniques to the content, structure, and usage of Web resources*" (Berendt et al., 2002). Such web mining techniques are also often used to support the creation of the semantic web (Berendt et al., 2002), as one part of the semantic web mining (Stumme et al., 2006), such as ontology learning, mapping and merging of ontologies, instance learning etc. However, this is not in focus of this paper. Instead we focus on knowledge discovery in the semantic web, which requires data mining techniques that can cope with the logical formalisms of the semantic web, to be more precise, relational data mining algorithms.

2.3 Relational Data Mining

Traditional data mining techniques can only handle data in a limited representation language, which was often propositional. This means that the data was transformed into a single table with an attribute-value structure. Instead, relational data mining builds upon the solid and expressive theoretical foundations of first-order logic (Raedt, 2008; Dzeroski and Lavrac, 2001; Knobbe, 2006; Dzeroski, 2003). Such relational algorithms are especially favourable in situations where learning problems involve multiple entities and relationships amongst them.

The research community has transformed the traditional algorithms for association rules, predictive modelling, clustering, statistical learning etc. to its first-order variants (Dzeroski and Lavrac, 2001). Such algorithms have been successfully applied in different scenarios such as genetics, molecules, social network analysis, as well as natural language processing (Dzeroski, 2001). However, while there has been detailed research on applying them on first-order logics, there has been only limited research on applying those techniques to the description logics (i.e. OWL-DL) of the semantic web.

Berendt et al. (2003), Berendt et al. (2002) and Stumme et al. (2006) provide a roadmap and initial starting point for this special emerging field of research. There are also already some application scenarios. Tresp et al. (2008) give an overview about different relational techniques and apply an infinite hidden relational model on friend-of-a-friend (foaf[5]) semantic data to recommend new friendships. Caragea et al. (2007) describe a relational bayesian classifier i.e. for the classification of computer science research papers in the bibliography domain. However, a real world evaluation is missing. Breaux and Reed (2005), Maedche and Zacharias (2002), Grimnes et al. (2008) and Fanizzi et al. (2008) present approaches for clustering entities with ontologies.

Despite of these promising approaches, the interest in this research topic is relatively small in comparison to the overall data mining research, especially in the context of business intelligence. Therefore we provide a scenario to motivate this field of research in the context of business information to outline the differences to the traditional knowledge discovery process. Furthermore, the utilization of meta-information based on description logics and relational data mining algorithms leads to changed requirements in this process, which we deduce from the research context as well as our scenario.

3 Scenario

In this section we describe an example of relational association rule analysis in the context of the semantic web, to be more precise, we aim to find frequent patterns in market basket data of an online shop. Furthermore, we assume that there is a data set that describes the interests of our customers, which are retrieved from a social network. This data set is similar to the foaf representation. We assume that there is a customer identifier (i.e. OpenID[6]) that is used in different portals, social networks etc. that can be used to identify persons in our data sets. An association rule analysis for market basket data is a standard analysis in this context (Han and Kamber, 2006). Therefore this is also an appropriate example to compare the propositional and relational approach in the knowledge discovery process to outline differences. Furthermore, we focus on the types of patterns that can be retrieved as well as their expressivity.

3.1 Data Cleaning, Integration, Selection and Transformation

The market basket data is stored in a simple relational database. The related Entity Relationship Diagram (ERD) is outlined in Figure 1. The corresponding tables as well as their data are shown below.

[5] http://xmlns.com/foaf/spec/

[6] http://openID.net

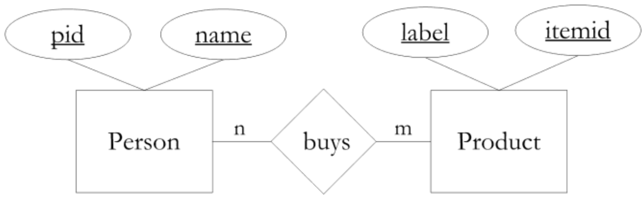

Figure 1 : ERD

itemid	label
Item1	Watersuit
Item2	Towel
Item3	Movie
Item4	Book

Table 1: Product

pid	name
P1	Maximilian
P2	Charlotte
P3	Stefan
P4	Kathrin
P5	Micheal

Table 2: Person

id	tid	itemid	pid
1	1	Item1	P1
2	1	Item2	P1
3	1	Item3	P1
4	2	Item2	P3
5	2	Item5	P3
6	3	Item3	P4
7	4	Item1	P5
8	4	Item2	P5
9	5	Item1	P6
10	6	Item1	P6
11	6	Item2	P6
12	7	Item1	P9
13	8	Item1	P9

Table 3: Person buys Product

The second data set is retrieved from a social network and contains background knowledge about interests and friends of agents. Below we outline the *ABox* of this data set.

Person(P1), Person(P2), Person(P3), Person(P4), Person(P5), Person(P6), Person(P7), Person(P8), Person(P9), Person(P10), Person(P11)

name(P1,“Maximilian“), name(P2,“Charlotte“), name(P3,“Stefan“), name(P4,“Kathrin“), name(P5,“Michael“), name(P6,“Martina“), name(P7,“Anja“), name(P8,“Kevin“), name(P9,“Marcel“), name(P10,“Moritz“), name(P11,“Bettina“)

Interest(I1), Interest(I2), Interest(I3), Interest(I4), Interest(I5)
label(I1,“Dolphins“), label(I2,“Waterski), label(I3,“Movies“), label(I4,“Comic), label(I5,“Dancing“)
hasInterest(P1,I1), hasInterest(P3,I2), hasInterest(P3,I4), hasInterest(P4,I3), hasInterest(P5,I5), hasInterest(P6,I5), hasInterest(P7,I1), hasInterest(P7,I2), hasInterest(P8,I1), hasInterest(P8,I2)

hasFriend(P1,P7), hasFriend(P1,P8), hasFriend(P5,P11), hasFriend(P6,P8), hasFriend(P9,P10)

In the standard propositional approach we are forced to transform the information into one table of transactions in form of an attribute-value structure (see Table 4). The modelling of the data is complicated and requires the introduction of several attributes. Of course it would be possible to analyse only the items of the transactions (white columns), but this would neglect the background information about the customers interests and their friends (grey columns).

Table 4: Transformed Data

tid	hasItem1	hasItem2	hasItem3	hasItem4	customer	hasInterest1	hasInterest2	hasFriend1	hasFriend2
T1	Item1	Item2	Item3		P1	Dolphins		P7	P8
T2	Item2	Item5			P3	Comic	Waterski		
T3	Item3				P4	Movies			
T4	Item1	Item2			P5	Dancing		P11	
T5	Item1				P6	Dancing		P8	
T6	Item1	Item2			P6	Dancing		P8	
T7	Item1	Item2			P9			P10	

Instead, the relational approach significantly simplifies the modelling of information. We define concepts for products and transactions as well as roles (properties) on them to create a small corporate ontology (a specific DL *TBox*). The instances of the database schema are transformed into the *ABox* of the DL based on the defined concepts and roles. The statements of this *ABox* are outlined below.

Product(Item1), Product(Item2), Product(Item3), Product(Item4) label(Item1,“Watersuit“), label(Item2,“Towel“), label(Item3,“Movie“), label(Item4,“Book“)
Transaction(T1), Transaction(T2), Transaction(T3), Transaction(T4), Transaction(T5), Transaction(T6), Transaction(T7)
hasItem(T1,Item1), hasItem(T1,Item2), hasItem(T1,Item3), hasItem(T2,Item2), hasItem(T2,Item5), hasItem(T3,Item3), hasItem(T4,Item1), hasItem(T4,Item2), hasItem(T5,Item1), hasItem(T6,Item1), hasItem(T6,Item2), hasItem(T7,Item1), hasItem(T7,Item2)
customer(T1,P1), customer(T2,P3), customer(T3,P4), customer(T4,P5), customer(T5,P6), customer(T6,P6), customer(T7,P9)

In OWL/RDF the information are represented as triples and thus can be directly automatically merged with the statements that provide the interest and friendship information, if the ontology concepts and roles are appropriately aligned. In this scenario, the complete *ABox* of the underlying OWL-DL ontology, which builds the basis for the relational association rule algorithm, is based on the merged *ABox* statements of both data sets.

Requirements

The data acquisition and transformation often requires 70%-80% of the time of the KDD process. In context of the semantic web this process should be automated partially, based on the semantic description of information, because the semantic web is a large scaling distributed system. In the scenario additional data about interests and friends of persons has been retrieved as additional background information, which is may be useful in the data mining step.

It seems therefore important to consider how software systems are able to (1) intelligently exploit relevant parts of the huge amount of available semantic information, (2) operate on a large-scale (d’Aquin et al., 2008; Berendt et al., 2002), (3) consider the high dynamic behaviour (Han and Kamber, 2006, p.628-629) of information (i.e. semantic web services) and (4) assess the quality of information.

Furthermore, the selection and transformation has to address semantic heterogeneities, if resources are not linked and/or ontologies are not aligned.

In addition, it is important to provide software architectures that allow a (5) transformation of company information from relational databases and/or data warehouses to the description logic formalism.

3.2 Data Mining, Interpretation and Evaluation

If we apply a standard association rule analysis, as described in Agrawal et al. (1993) on attributes of the white part of Table 4, we get i.e. the following item sets with minimum frequency 0.4.

{Item1,Item2}	FREQUENCY: 4/7
{Item1}	FREQUENCY: 5/7
{Item2}	FREQUENCY: 5/7
Rule: {Item1}→{Item2}	CONFIDENCE: 0.8
Rule: {Item2}→{Item1}	CONFIDENCE: 0.8

Thus we know that it is likely (confidence 0.8) that a customer will buy "Item2", if a customer buys "Item1" and vice versa. If we also consider the background knowledge from the transformed data set (grey part of Table 4), we get also the following item sets with minimum frequency 0.4.

{Item1,Dancing}	FREQUENCY: 3/7
Rule: {Item1}→{Dancing}	CONFIDENCE: 3/5
Rule: {Dancing}→{Item1}	CONFIDENCE: 1.0
{Item1,P8}	FREQUENCY: 3/7
Rule: {Item1}→{P8}	CONFIDENCE: 3/5
Rule: {P8}→{Item1}	CONFIDENCE: 1.0

However, if we use relational data mining, we can consider the complete available background knowledge. In relational data mining, we know search for relational item sets in the *ABox*. A relational item set (Dzeroski and Lavrac, 2001) is a kind of query on the knowledge base *(KB)*. A query is denoted by a "?". A query succeeds if there is an answer (i.e. ?-Product(x) has answers Item1,Item2,…,Item4) in the knowledge base. We have to specify the base concept of interest, which in our scenario is the transaction. We can now get relational item sets as follows (X,Y,Z,P,F are different query variables). The queries are executed on the merged *ABox* statements of our knowledge base described above.

?- Transaction(X)	FREQ.: 1.0
?- Transaction(X), hasItem(X,Y)	FREQ.: 1.0
?- Transaction(X), hasItem(X,Y), hasItem(X,Z)	FREQ.: 5/7
?- Transaction(X), hasItem(X,Y), customer(X,P)	FREQ.: 1.0

?- Transaction(X), customer(X,P), hasFriend(P,F), hasInterest(F,"Waterski") FREQ: 5/7

?- Transaction(X), hasItem(X,"Item1"), hasItem(X,"Item2"),
customer(X,P),hasFriend(P,F), hasInterest(F,"Waterski") FREQ.: 4/7

Rule: Transaction(X), customer(X,P), hasFriend(P,F), hasInterest(F,"Waterski") → hasItem(X,"Item1"), hasItem(X,"Item2") CONFID.: 0.8

We know from this relational rule (with a confidence of 0.8) that a customer will buy "Item1" and "Item2", if he has a friend that is interested in "Waterski". This is a much more expressive rule than in the propositional approach, which considers the given background knowledge.

Algorithms that create the relevant item sets can be found in Agrawal et al. (1993) and for relational item sets in Dzeroski and Lavrac (2001). However, in this paper we focus on the types of patterns. Even in this simple example, relational data mining improves the expressivity of the rules.

Requirements

The requirements for the data mining step increase, because the distributed nature of the information and the large amount of information requires very (7) efficient algorithms.[7] Furthermore, these algorithms need fast reasoning support to execute queries on the knowledge base.

The selection of an appropriate data mining algorithm is a task of the decision maker or analyzing person based on the specific goal. However, the combination of relational data mining, semantic web, logic etc. is very complex and it seems to be not likely that managers have a great interest to understand all details (Fayyad, 2007). It is therefore important to (8) perform research on managerial utilization of the proposed technology. This requires research on appropriate user interfaces and application architectures semantic web based data mining.

The consideration of relational patterns increases the amount of possible useful ones. Thus it becomes more and more important to perform research on how to extract useful patterns for a decision maker.

Furthermore, it is important to specify evaluation methodologies for data mining algorithms on relational data. This requires especially the generation of artificial test sets that enable a detailed evaluation. Real world test sets are not appropriate as a standalone evaluation, since the real distribution of the data is unknown.

4 Conclusion and Future Research

In this paper we have motivated the utilization of relational data mining algorithms in context of the semantic web. In our specific scenario, we have outlined important differences to the traditional knowledge discovery process. Especially the selection, modelling and transformation steps are different to the standard approach. As described in the scenario, data from data warehouses, databases as well as the

[7] The LinkingOpenDataProject (http://www.linkeddata.org) provides currently semantic information based on more than one billion triples.

web has to be transformed into the logical formalisms of the semantic web (see Figure 2). Then the knowledge discovery process selects parts of this semantic data to fully utilize information and background information derived from a variety of sources. In accordance to our scenario as well as the related work this can improve the expressivity and amount of patterns that can be found in the subsequent relational data mining step.

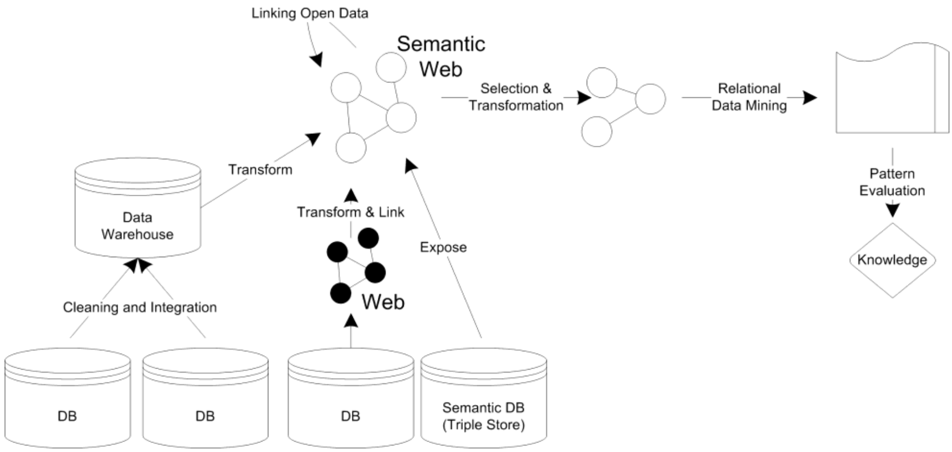

Figure 2: Knowledge Discovery in the Semantic Web

There are existing application scenarios of relational data mining especially in biology, but it remains open how knowledge discovery in context of the semantic web can be used in enterprises (i.e. risk management, competitive analysis etc.). In our ongoing and future work, we aim to develop a prototype that addresses the stated requirements. Furthermore, we aim to outline and evaluate specific company scenarios and decision processes that benefit from these techniques.

References

Agrawal R, Imielinski T, Swami A (1993) Mining Association Rules between Sets of Items in Large Databases. SIGMOD Conference:207-216.

Baader F, Calvanese D, McGuinness DL, Nardi D, Patel-Schneider PF (2003). The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge, Cambridge.

Berendt B, Hotho A, Mladenic D, van Someren M, Spiliopoulou M, Stumme G (2003). A roadmap for web mining: From web to semantic web. In: EMWF 2003, Cavtat-Dubrovnik, Croatia, September 22, 2003, Springer.

Berendt B, Hotho A, Stumme G (2002). Towards Semantic Web Mining. In: ISWC 2002, First International Semantic Web Conference, Sardinia, Italy, June 9-12, 2002, Springer.

Berners-Lee T, Hendler J, Lassila O (2001) The semantic web. Scientific American: 29–37.

Breaux, TD, Reed JW (2005) Using ontology in hierarchical information clustering. In: HICSS, 3-6 January 2005, Big Island, HI, USA. IEEE Computer Society.

Brickley D, Guha R, McBride B (2004) RDF Vocabulary Description Language 1.0: RDF Schema. http://www.w3.org/TR/rdf-schema/

Caragea D, Bao J, Honavar V (2007) Learning relational bayesian classifiers on the semantic web. In: IJCAI Workshop on Semantic Web for Collaborative Knowledge Acquisition (SWeCKa).

d'Aquin M, Motta E, Sabou M, Angeletou S, Gridinoc L, Lopez V, Guidi D (2008) Towards a New Generation of Semantic Web Applications. IEEE Intelligent Systems, 23(3):20–28.

Dzeroski S (2003) Multi-relational Data Mining: An Introduction. SIGKDD Explorations, 5(1):1–16.

Dzeroski S and Lavrac N (2001). Relational Data Mining. Springer, Berlin.

Edelmann W (2000). Lernpsychologie. Kösel, Verlagsgruppe Beltz.

Elomaa T, Mannila H, Toivonen H (2002) Principles of Data Mining and Knowledge Discovery. In: PKDD, Helsinki, Finland, August 19-23, 2002. LNCS, Springer, Berlin.

Fanizzi N, d'Amato C, Esposito F (2008) Conceptual clustering and its application to concept drift and novelty detection. In: ESWC 2008, Tenerife.

Fayyad UM (2007) From mining the web to inventing the new sciences underlying the internet. In: Berkhin P, Caruana R, and Wu X, editors, KDD: 2–3. ACM.

Fayyad UM, Piatetsky-Shapiro G, Smyth P (1996) From Data Mining to Knowledge Discovery in Databases. AI Magazine, 17(3):37–54.

Grimnes GA, Edwards P, Preece AD (2008) Instance based clustering of semantic web resources. In: ESWC 2008, Tenerife, Canary Islands, Spain, June 1-5, 2008, LNCS:303–317. Springer.

Gruber TR (1993) Towards principles for the design of ontologies used for knowledge sharing. In: Formal Ontology in Conceptual Analysis and Knowledge Representation, Kluwer.

Han J, Kamber M (2006) Data Mining: Concepts and Techniques, Second Edition, Morgan Kaufmann.

Klyne G, Carroll JJ, McBride B (2004) RDF Primer. http://www.w3.org/TR/rdf-primer/

Knobbe A (2006) Multi-Relational Data Mining. Frontiers in Artificial Intelligence and Applications (145). IOS Press.

Krcmar H (2004). Informationsmanagement (German Edition). Springer, Berlin.

Maedche A, Motik B, Stojanovic L (2003) Managing multiple and distributed ontologies on the Semantic Web. VLDB J., 12(4):286–302.

Maedche A, Zacharias V (2002) Clustering Ontology-Based Metadata in the Semantic Web. In: Elomaa et al., 2002: 348–360.

Müller-Merbach H (2004) Knowledge is more than information. Knowledge Management Research & Practice, 2(1):61–62.

Raedt LD (2008) Logical and Relational Learning (Cognitive Technologies). Springer, Berlin.

Rowley J (2007) The wisdom hierarchy: representations of the dikw hierarchy. Journal of Information Science, 33(2):163–180.

Singh MP, Huhns MN (2005). Service-Oriented Computing: Semantics, Processes, Agents. Wiley.

Smith MK, Welty C, McGuinness DL (2004) OWL Web Ontology Language Guide. http://www.w3.org/TR/owl-guide/

Stumme G, Hotho A, Berendt B (2006) Semantic Web Mining -State of the Art and Future Directions. Journal of Web Semantics, 4(2):124–143.

Tan PN, Steinbach M, Kumar V (2005) Introduction to Data Mining. Addison Wesley.

Tresp V, Bundschus M, Rettinger A, Huang Y (2008) Towards Machine Learning on the Semantic Web. In: ISWC International Workshops, URSW 2005-2007, LNCS:282–314. Springer.