Knowledge Fusion for Financial Advisory Applications

Tadeusz P. Dobrowiecki, György Strausz, Tamás Mészáros

Department of Measurement and Information Systems, Budapest University of Technology and Economics (BUTE), Mugyetem rkp. 9., Budapest, Hungary, H-1521, E-mail: dobrowiecki@mit.bme.hu[?]

ABSTRACT: In the paper we present a design of a decision support system for the fully automated controlled retrieval of fiscal information, which is being continuously collected from a variety of sources, mainly from the global computer network. This system extends the idea of data mining to the notion of knowledge retrieval. The collected information is processed in a knowledge-intensive way using knowledge models of the application domain and can subsequently be used to support various decision-making processes.

1. Introduction

Decisions in finances become more and more tedious and effective solutions cannot be reached without considering all of the essential influencing circumstances. Development of the explicit models of the target relations is hard, despite the huge amount of the available information.

This huge amount of data stored in information systems can be used to extract direct information about the clients and the transactions, with the aim of generating a competitive advantage and an improved quality of service. Successful applications - e.g. automatic extraction of information from data sources for the computer-based or human-based decision support - show that even without considering the meaning of the stored data, useful information, e.g. trends or anomalies, can be determined. *Data mining* deals just with such extraction of implicit, previously unknown, and potentially useful information and with using it for crucial business decisions [1, 2]. It uses machine learning, statistics and visualisation techniques to the discovery and the presentation of knowledge in a form that is easily comprehensible to the humans. Data mining deals primarily with structured and well-defined data sources, especially with large relational databases.

Information retrieval means collecting information from unstructured text documents (like books, papers, and electronic documents) [3]. In the 90s the easiness experienced in the Internet publishing resulted in a greatly increased volume of documents stored in global computer networks. To find the right piece of information in this rapidly growing and unstructured data storage, new kinds of systems were developed, like e.g. Internet catalogues and search systems. Automatic indexing systems attempt to cope with vast amount of data and to build structured index schemes. Internet search engines use them then to find the required information. There are also manually edited index schemes to be browsed by the humans. These systems will become increasingly powerful by integrating into their function the knowledge about the required information.

In the today's information retrieval and data mining the meaning of the information which is looked for is not used. Any related knowledge, however, can vastly increase the effectiveness of the retrieval process. *Knowledge retrieval* represents a significant advancement over simple search engines and conventional information retrieval methods.

2. Information & Knowledge Fusion (IKF)

[?]This research was supported by the Hungarian NSF (OTKA) Grant Nr. T-030586, and belongs to the IKMA IKF-H prolog research.

Efficient knowledge retrieval is being investigated within the framework of the Information and Knowledge Fusion EUREKA Applied Research Project. Our main goal is to analyse, design and implement new Intelligent Knowledge Warehousing environment, which would allow for an advanced knowledge management in various application domains (e.g. banking, revenue service, insurance, legal information, training & education, health care, etc.) [4].

The Hungarian IKF (IKF-H) consortium develops knowledge-based information retrieval system for financial corporations and banks [5]. The system will retrieve topical information from various sources (Internet, Intranet, data warehouses, corporation databases, expert human resources, etc.), and provide the information to the users in a structured way.

Large banks experience difficulties when collecting and evaluating information about their clients. Conventional information sources in such application areas provide only limited and frequently unreliable information. It is also hard to evaluate such data in an optimal way. Despite the fact that nowadays quite a variety of information sources provides relevant data about the clients of the banks, the human evaluation has limited depth, scope and reliability. Only very few sources are considered for gathering information and the human resources are in all respect limited.

To help this situation the project formulated a number of objectives, namely:

(1) to determine suitable methods for describing semantic information, that can be used to enhance the retrieval efficiency in financial information retrieval;

(2) to evaluate the state of the art of knowledge representation and reasoning methods, and to integrate these techniques into the retrieval process; and

(3) to build knowledge-based information retrieval system.

Such system can be used as an electronic alternative to the conventional data collection and analysis tools, which perform the client authorisation. It can provide new services, greater depth of the analysis, and higher reliability in the authorisation process. In particular it can:

?? **Continuously monitor the clients** by automatically updating the related knowledge (about the clients) and updating the authorisation data. The system frequently checks the available information sources and learns about changes related to the tracked clients. The system is able to analyse the new data and report the new circumstances to the customer (bank), if necessary.

?? **Select the relevant data and providing reliability information** as well**.** With the system the user can select from a wide range of relevant information, as the system collects data from several sources. The system also provides an estimated measure on the reliability of the data.

3. Abstract Model of the Problem Domain

For further considerations it must be noted, that at a properly abstracted level all of the applications mentioned above can be modelled as a formation and a purposeful interaction of three *information environments*. *Target Environment* is that fragment of business life, where the targeted (monitored) corporations and clients do exist. *Information Cumulating Environment* comprises all forms and media, which cumulate information about the targets. In our case it is Internet, various Intranets, corporation databases, published resources, financial experts, etc. *Information Utilising Environment* means e.g. the staff of a bank, at various level of management, see Fig.1.

Consequently a structured raw and derived information (i.e. knowledge) is flowing between the Target and the Information Cumulating Environments and from the knowledge-intensive Information Cumulating to the Information Utilising Environment, where specific knowledge is needed to fulfil specific goals. In the following, when convenient, we will address these three environments as the "Client", the "Web", and the "Bank" respectively.

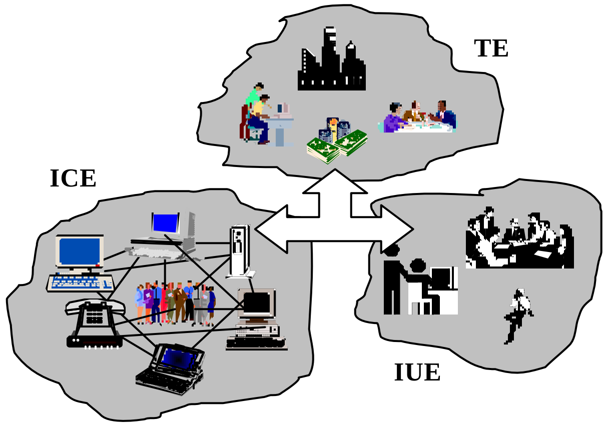

Fig. 1. The three interacting environments.

Target Environment (TE, "Client") is the physical source of knowledge. It comprises objects, phenomena, relations, etc. whose particular properties (parameters, relations, and state variables) are of interest for both Information Cumulating ("Web") and Information Utilising Environments ("Bank"). Target objects are usually interrelated, i.e. part of the knowledge is common to whole subsets or structures of domain objects.

Information Cumulating Environment (ICE, "Web") is coupled to the Target Environment by knowledge acquisition process of various sorts. Knowledge embedded in the "Web" is thus to an extent a veritable model of the "Client". It is however important to note that such knowledge is heavily distributed within the "Web" and that the application (i.e. "Bank") has no control over the extent and the timeliness of the acquisition process. Consequently the "Bank" has no control over what knowledge is stored and where.

Information Utilising Environment (IUE, "Bank") represents active and possibly intelligent entities that require particular knowledge about the objects from within the "Client" to achieve their specific goals within the "Bank". An interesting feature adding to the complexity is that the "Web" and the "Bank" can be related, even overlapping. This will be true if e.g. the "Web" involves knowledge cumulated in human (expert) resources.

4. Knowledge Level

Analysis of the knowledge involved in the knowledge-intensive problems, so called knowledge-level analysis, is recommended to grasp the essence of the difficulties before the involvement in the details of the representation and the implementation could screen some essential problems.

The knowledge stored in the "Web" and the "Bank" is the derivative of the knowledge characteristics of the "Client". Efficient information retrieval is difficult, because the "Client" reveals a great dose of complexity (due to the complexity of the individual objects and to the rich structure of their relations).

Why the target objects seem complex? Generally speaking they seem so, because they are distributed in the environment, admit multiple perspectives and abstractions, evolve in time, and typically represent uncertain and missing information.

Distribution within the environment means that at least a part of target objects appears distributed from the point of view of the knowledge acquisition process run by the "Web". Consequently their characteristics will be stored in the "Web" in different forms and at different locations.

An inherent characteristic of the target objects in the considered domains is their complexity. This ***complexity*** means that an object can be described, understood, simulated, etc. only when utilising a large number of parameters or state variables. Even if, temporarily, it could be characterised completely by a smaller number of parameters, such parameters are usually computed, and not measured directly from the whole spectrum of the object-related parameters. Complexity means also that the parameters are related through structures and those structures are also characteristics to the particular classes of objects.

Objects in the chosen application areas are in themselves complex systems comprising diverse phenomena. Consequently, when monitoring them, the necessary selection and limitation in observations will produce ***different perspectives*** of the very same object. Due to the above stated reasons a particular object can be monitored, or viewed at various levels of abstraction or detail. It is plausible to assume that ***different levels of abstraction*** will be used in the same time with respect to the same object by different entities in the "Bank".

Target objects in "Client" are necessarily related and connected by the flow of certain common commodities (e.g. manufactured materials, financial means), making "Client" a ***highly structured*** and organised universe. Although this increases the complexity of the maintenance of any kind of model, it also presents an opportunity to introduce more involved knowledge deriving mechanisms, e.g. ***logical inference*** or ***inheritance of properties***.

Although certain objects in the "Client" are static, the majority of the target objects is dynamic by nature. Consequently their ***properties,*** even ***certain elements of their structure will change in time***. It does not necessarily mean a problem for the "Web"; it will be, however, almost surely reflected in the dynamic character of the goals, making the transitions of knowledge between the "Web" and the "Bank" difficult. As an additional twist in difficulty and a direct consequence of the distributed approach, multiple perspectives and abstraction levels, the knowledge collected within the "Web" will necessarily *be* ***fractional, missing and uncertain***. This will add to the burden of how to design the transition of knowledge between the "Web" and the "Bank" in such a way as not to jeopardise the fulfilment of the "Bank" goals.

5. Mediating the Fusion and the Dissemination

Our primary and basic problem to face is that the knowledge acquisition between "Client" and "Web", and its storage in the "Web" are not specific to the goals found in the "Bank". The character of the knowledge acquisition reflects rather the physical and technological possibilities of the "Web" or some unrelated goals pending within that environment. Consequently the form, the quality, and the validity of information are questionable from the users' point of view. No single entity in the "Web" can be directly addressed and pointedly requested for information.

What is required to solve this problem is a kind of means-ends analysis where the "Web's" means must meet the "Bank's" ends. Due to the dynamic character of the "Client" and the continuous shift in the "Bank's" goals, such analysis should be conducted on the round-the-clock basis and implemented that way by the suitable systemic means.

Specific goals require more condensed information than that available in the "Web's" sources. On the other hand when the user goals vary, so does vary the character of the requested information, even about the same "Client's" objects. In consequence the "Web" and the "Bank" must be interfaced by a mediating system embedded in both, which can accept as input the "Bank's" goals and which to provide answers should be familiar with the "Web's" possibilities and limitations, see Fig. 2.

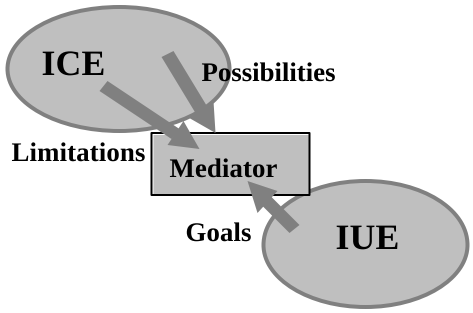

Fig. 2. Mediating system.

The components of the required system follow logically from its function. Such system must help to find in the "Web" the information essential for the "Bank's" goals, however, for typically recurrent and related goals producing the information by the repeated request to the "Web" would mean an unjustified and badly organised spending of system resources. Consequently a Knowledge Repository (KR) is needed, which reflects a portion of the "Client" and would puffer the information to avoid spending too much of the system resources on extracting from the "Web" knowledge needed for recurrent and related retrieval requests.

Repository must be up-dated with intensity and agility defined by the "hunger" for the information in the "Bank" and with the information accumulated in the meantime in the "Web" The question is therefore where from and in what form seek the information? To handle this problem the Information Retrieval Subsystem (IRS) is required, which will perform knowledge sensitive and knowledge intensive information retrieval. The IRS should know models of the "Client" (essentially the structure of the KR), should know the model of the "Web", and should perform search for fresh information.

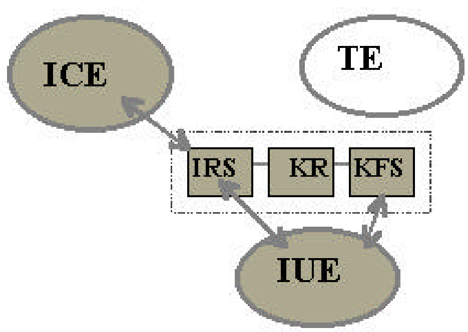

Fig. 3. Internal structure of the mediating component.

To achieve a truly flexible solution the KR should keep more diverse information than what is demanded by the temporary needs of the users. When the goals shift (which is normal in the considered application areas), the requirements with respect to the type and the volume of information about the "Client" will also change. Nevertheless, the KR must continuously provide it. Within reasonable limits such information should be already available in the KR, but it should be served in the proper format and extent to the designated entities in the "Bank". To translate the new goals into the new requests toward the KR we need a new system components – Knowledge Fusion Subsystem (KFS), see Fig. 3.

The Knowledge Repository is composed primarily from the properties of the physical objects from the "Client", filled-in by the IRS. To permit goals and queries related to groups and hierarchies of objects, the design of the Knowledge Repository should also involve a suitable hierarchical model of the "Client".

Moreover KR should be able to serve information not necessarily introduced directly by the IRS component (i.e. relations, more abstract fact, trends, etc.). Two problems are to be considered here. Firstly, the updating of the KR should be quicker than the expected highest sampling frequency of the trends as asked in the queries. It would be possible, if the IRS were familiar with the "Client" goals.

Secondly the KR must be equipped with the suitable knowledge derivation (inference) mechanism, making it rather a knowledge-based (expert) system and not a traditional database. This would allow the users to formulate information retrieval goals at various levels of abstraction and correlating multiple objects from the "Client".

Although the KR is a well-defined functional component, physically it can be organised as a distributed system, to facilitate the implementation. The same applies of course to other system components. Whether the parts of the mediating system will be localised or distributed, it ultimately depends upon the characteristics of the problem domain. Due to the distributed character of the "Web" and the "Bank" localised within the corporation limits, a distributed (and possibly highly autonomous, agent-based) IRS and localised KFS and KR components are the most probably implementation options.

In the following we review the interactions between the environments and the system components and draw necessary consequences completing the generic architecture of the mediating system.

The activity of the IRS is composed from the following three interaction:

?? IRS ? "Web" - looking for sources of information covering the goals of the "Bank",

?? IRS ? KR - storing the retrieved information, and

?? IRS ? "Bank" - accepting hints and prescriptions about what, how, and where to fetch.

The first two interactions are natural, but the third one requires more deliberation. Information richness of the "Web" and finite system resources (e.g. limited memory in KR, limited access time, etc.) do not permit to follow all the possible paths in the "Web" and all the possible objects in the "Client" in their full complexity. The IRS must filter what it is looking for. Such filtering, however, must be somehow defined by the "Bank" to maintain a meaningful KR and an acceptable level of services.

Definition of the filtering should belong to the management level within the "Bank". The "Web" may be considerably diversified, covering distant network resources, specialised databases, and human resources reachable by telephone, fax, e-mail, etc. It implies that depending upon the application, the characteristics of the "Web", the goals of the "Bank", etc. the acquisition and the refreshing process can be indeed diversified and the IRS will require help from qualified operators, experts in the overall goals and strategies of the "Bank".

Similarly to the IRS the Knowledge Fusion Subsystem interacts with the KR and the "Bank". The Repository, accordingly to what was said earlier, has to “know” more than the particular needs of the particular users in order to serve the diversified goals of the whole "Bank". Therefore, the KFS component must be told what to fetch in the KR and in what form to present it to the ordering entity in the "Bank".

In the following let us shortly review various support functions provided by the system. From the design point of view the "Bank" is composed from end-users and managers. End-users ask simple queries and accept data. Managers, in addition, shape the overall flow of information.

The information flow at the end-user level (i.e. the level of the individual goals in the "Bank") is governed by the Report Profile - a concise representation of the user's goals - stating which objects, parameters, etc. should be fetched from the KR and transformed into actual personalised reports, Fig. 4. The KFS finds the required information (continuously up-dated by the IRS), and reports it to the users in the "Bank" with automatically generated periodic and one-shoot summaries. The KFS works with a number of report profiles.

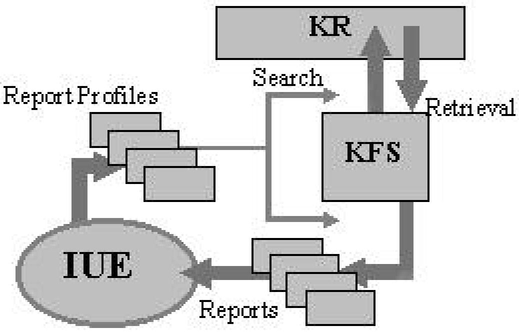

Fig. 4. Information flow related to the Knowledge Fusion Subsystem.

Similarly to the user defined Report Profiles managers must define with suitable Search Profiles what, how, and where to search. Search Profiles modify the searching strategy within the "Web" and the updating strategy within the KR, Fig. 5. Managers can also add Report Profiles to define reports to be distributed all over the "Bank" (e.g. corporation news, daily or hourly exchange rates, etc.).

The key part of the system is the internal Knowledge Base that models the application area and in particular the application task. The model building process is difficult and requires expertise in both the internal structure of the IKF-H system and in the specific application field. Therefore only system managers can use development and management services. These tools are hidden at the user level. Main tasks of the management services are the followings:

(1) *Design and modification of the Knowledge Base*. The Knowledge Base contains general and application specific components. The main purpose of this decomposition is to allow the system manager to design the domain specific information. System manager also can analyse and modify any part of the internal system data.
(2) *Defining and modifying search and evaluation strategies*. A set of information search and processing functions is available in the system. System manager can select methods appropriate to the tasks and can also modify them by setting the parameters of the methods. This customisation includes the specification of data sources, and retrieval and processing methods as well. The reliability of these sources can also be described.
(3) *System performance analysis*. In order to be able to accomplish the above mentioned tuning functions system manager should be able to analyse the behaviour of the system. A wide variety of analysing functions are necessary to accomplish the general utilisation on the system.
(4) *Monitoring, and supervising*. Monitoring and supervising is necessary in every complex system, but this task is especially critical for systems with adaptive nature.
(5) *Input form specification*. One of the important tasks of system manager is to create application specific input forms for the user. Designing input forms means not only generating fields to fill-in, but also describing the relations between the fields and the internal representation.
(6) *Output report structure definition*. System manager can determine what part of the information should be recallable from the Knowledge Base. System manager can also define what part of the Knowledge Base is reachable for the end-users and what part should remain hidden for them. Pre-defined report forms can also be created and customised. The user can select these forms during the report creation.

User level services provide the interface for the end-users to define their requests and to read out the outputed results. The basic strategy applied here is that the end-user has only limited access to the system services. End-users can only reach information relevant to their working topics and are offered basically three types of services:

(1) *Generating queries using pre-defined input forms*. End-users can input information into the Knowledge Base and also pose queries to the system through pre-defined forms. End-users have

limited possibility to configure the forms from the list of entities defined earlier by the system manager.

(2) *Generating output reports by using and customising pre-defined report forms*. Users have similar possibilities for defining output report forms like in case of the input forms. This way end-users can produce dense reports on a given topic.

(3) *Asking for the information about the results of a query*. An important service is to provide explanation about the output data and also how the output data was generated. The system offers also an interface service to the users to learn about the internal models of the system.

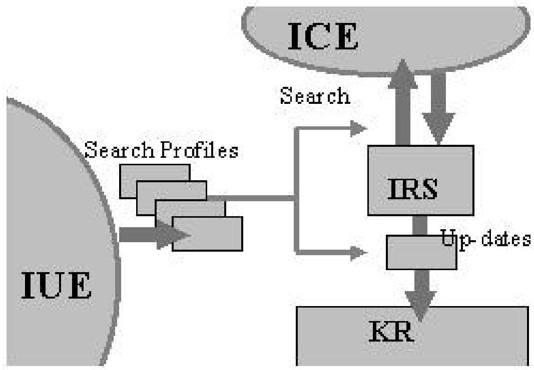

Fig. 5. Information flow related to the Information Retrieval Subsystem

6. System Integration Issues

"Web" contains a very diverse set of information repositories starting from relation databases, Internet resources up to the interfaces to various human data resources. In order to gather all these data the IRS should be equipped with several kinds of properly customised retrieval methods. We can categorise them as follows:

(1) *Data queries*. Retrieving data from structured and well-defined data sources (like databases).

(2) *Data mining*. An automated approach to exhaustively explore and determine complex relationships in very large databases [1, 2].

(3) *Information retrieval (IR)*. IR deals with the retrieval process from unstructured data sources, e.g. the Internet.

"Bank" requires methods for human-computer interaction (HCI) - graphical program interfaces, models of human roles and behaviours, etc. - that can be used to enhance the effectiveness of the system. The analysis of the "Client" leads to the selection of the knowledge-based modelling techniques and reasoning methods. From the analysis presented in the previous chapters it clear that the system must deal with structured and inter-related information with reliability measures.

The distributed nature of the problem suggests a generic multi-agent architecture [6, 7], where the main components, and the parts of those are described as individual agents that co-operate or compete to fulfil a their own goals and thus the overall system aims. An intelligent agent is a software system that operates autonomously to fulfil locally specified goals. The main characteristics of these systems are reactivity (it senses its environments and makes actions based on the perceptions), proactivity (it works to fulfil its goals), social ability (it is able to communicate with other agents and humans), and persistency (it continuously maintains an internal state). Other secondary characteristics may include being veritable, adaptive, robust, rational, and mobile.

There are emerging standards in this field, like the ARPA Knowledge Sharing Effort (KSE), the OMG

Mobile Agent Facility (MAF), and the Foundation for Intelligent Physical Agents (FIPA) [8]. Multi-agent systems (MAS) contain autonomous agents that communicate with each other to solve common tasks. This requires a formalised communication language that can be implemented by the creators of the individual agents. The KQML (Knowledge Query and Manipulating Language) is and example of such a language [9]. There are two basic types of MAS: co-operating and competitive ones. Agents in co-operating systems share results and information to solve the common problem. Competitive MAS contains individual agents that operate on the same problem, however better problem solvers gain more resources (rewards), therefore achieving better position for the solution of future problems. Both ideas are applicable from the IKF-H perspective.

One of the key components of the IKF-H system is the Information Retrieval Subsystem. It is responsible for gathering the information from the "Web" according to the expectations of the "Bank". Recently much of R&D effort is being made to reduce, filter, organise, and present the mass of data and information available from the computer networks, and to integrate them into dynamic knowledge models. Agent systems are used here for:

(1) discovering suitable information sources,
(2) gathering and filtering information,
(3) up-dating and refining the knowledge model.

These systems are based on the following technological methodologies:

(1) intelligent, agent-based information retrieval,
(2) text document analysis and natural language processing,
(3) knowledge-based information exploitation,
(4) domain specific information structuring and reasoning,

Different agents can be designed and customised for different information sources. The multi-agent architecture can also be used to increase the reliability of the collected data. Co-operating agents will deal with the different information sources, and competing agents can be utilised to enhance the retrieval effectiveness. The agent methodology can be also applied during the interaction with the user. The complexity of the IKF-H system and the problem domain can be hidden from the user using specialised agents. Simple interface agents have tasks like helping the user in filling-in forms, or automatically correcting errors in user inputs. More sophisticated agents can be used as application helpers, explaining and teaching the usage of the system.

Regarding the details of the Knowledge Repository two leading concepts are the knowledge organised around the suitable domain ontology and the XML-based document retrieval, mapping and storage. An ontology is an explicit (possibly formal) specification of the names for referring to the objects in the application and the (logical) statements that describe what these objects are, and how they are related or not to each other [10, 16-17, 20].

Ontology therefore provides a vocabulary for representing and communicating knowledge that can exist for an agent or a community of agents, for the purpose of enabling knowledge sharing and reuse. The so-called ontological commitment means an agreement to use a vocabulary (i.e. for queries) in a way that is consistent with the theory. Research in ontology is one of the most far-reaching issues in the now-a-days artificial intelligence [18].

The crucial role of well-defined ontology has already been recognised. IEEE Computer Society pioneered a Standard Ontology Study Group [11]. A strong research group related to the IKF project is tackling the question of the meta-organisation of the ontological hierarchies [19], finally fairly recently a number of ontology based enterprise models has been developed [12-15]. All these developments serve as a basis for the development of the suitable Hungarian enterprise ontology and knowledge base.

After successful attempts to standardise on the hardware and software XML is an attempt to standardise on the information, to make efficient search, storage and retrieval possible. XML is so called mark-up language, i.e. a set of mark-up conventions to encode the meaning of the text. XML is extensible and various concrete languages can be created for particular domains, with the automatic translation support [21-23].

From the IKF point of view the most interesting developments are: XML-Based Ontology Exchange Language [29], Financial Products Mark-up Language [27], Artificial Intelligence Mark-up Language [32], DARPA Agent Mark-up Language [24, 28], Extensible Financial Reporting Mark-up Language [30], Trading Partner Agreement Mark-up Language [31], Extensible Business Reporting Language [26], and the like [25].

Here similarly the main R&D task will be the shaping the emerging tools to the in-country requirements of the "Bank" environment.

7. Summary

The collection of information form a vast number of information sources available through the global computer networks and a knowledge intensive processing (reduction) of such information with the purpose to back up management decisions and evaluations opens new dimension in the effective financial management.

Designing information systems, which providing such support is already possible now-a-days, if a number of emerging software, information processing, and system integrating technologies is used and further developed to the needs of the application. Designed approaches, tools and methods are portable to numerous domains, where the abundance of information and difficulty of making decisions made the introduction of the automated information system questionable until now.

References

[1] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, R. Uthurusamy (Eds), “Advances in Knowledge Discovery and Data Mining”, Mit Press, 1996

[2] Data Mining Paper Collection, http://www.ececs.uc.edu/~mazlack/papers.dbm.html

[3] Korfhage, R., “Information storage and retrieval”, Wiley Computer Publishing, 1997

[4] EUREKA PROJECT “IKF - Information and Knowledge Fusion”, March 2000.

[5] “Információ és Tudás Tárház (Information and Knowledge Warehouse)”, IKTA/3, 99102813, Vol. 2., Work Plan and Cost Statements, May 2000.

[6] J. M. Broadshaw, “Software Agents”, The MIT Press, 1997

[7] AI on the Web, http://www.cs.berkeley.edu/~russell/ai.html

[8] Y. Labrou, T. Finin, and Yun Peng, “Agent Communication Languages: The Current Landscape”, IEEE Intelligent Systems, March/April 1999, pp. 45-52

[9] UMBC KQML Web, http://www.cs.umbc.edu/kqml/

[10] J. F. Sowa, “Knowledge Representation: Logical, Philosophical, and Computational Foundations”, Brooks Cole Publishing Co., Pacific Grove, 1999

[11] Standard Upper Ontology, IEEE Study Group, IEEE Computer Society, Standards Activity Board, June 2000, wysiwyg://634/http://ltsc.ieee.org/suo/index.html

[12] The Enterprise Project: The Enterprise Ontology, 1998, http://www.aiai.ed.ac.uk/project/enterprise/enterprise/ontology.html

[13] J.Stader, “Results of the Enterprise Project”, in Proceedings of Expert Systems '96, the 16th Annual Conference of the British Computer Society Specialist Group on Expert Systems, Cambridge, UK, December 1996.

[14] M.S. Fox, J.F. Chionglo, and F.G. Fadel, “A Common-Sense Model of the Enterprise”, Proceedings of the 2nd Industrial Engineering Research Conference, 1993, pp. 425-429, Norcross GA: Institute for Industrial Engineers

[15] US Taxonomies, US GAAP C&I Taxonomy 00-04-04, http://http://www.xbrl.org/US/default.htm

[16] N. Guarino, "The Ontological Level", 16th Wittgenstein Symposium, Kirchberg, Austria, 1993

[17] John Sowa's web site devoted to knowlege representation and related topics of logic, ontology, and computer systems, http://www.bestweb.net/~sowa/direct/index.htm

[18] Some Ongoing KBS/Ontology Projects and Groups, http://www.cs.utexas.edu/users/mfkb/related.html

[19] N. Guarino and Ch. Welty, "A Formal Ontology of Properties", LADSEB/CNR Technical Report 01/2000, http://www.ladseb.pd.cnr.it/infor/ontology/Papers/OntologyPapers.html

[20] N. Guarino and Ch. Welty, "Towards a methodology for ontology-based model engineering", in, Bezivin, J. and Ernst, J., eds, Proceedings of the ECOOP-2000 Workshop on Model Engineering. June, 2000.

[21] P. Prescod and Ch. F. Goldfarb, "The XML Handbook - 2nd Edition", Prentice Hall, 1999

[22] The small XmeLler, http://www.sciences.univ-nantes.fr/info/lrsg/sNets/XML.html

[23] Parand Tony Daruger, "Tutorial: XML and scripting languages", Feb 2000, http://www-4.ibm.com/software/developer/library/xml-perl/

[24] B.N. Grosof and Y. Labrou, "An Approach to using XML and a Rule-based Content Language with an Agent Communication Language", IBM Research Report, RC 21491 (96965) 28 May 1999, Computer Science, http://www.research.ibm.com/

[25] XML: Proposed Applications and Industry Initiatives, http://www.oasis-open.org/cover/xml.html#applications

[26] Extensible Business Reporting Language (XBRL), http://www.oasis-open.org/cover/xbrl.html

[27] Financial Products Mark-up Language (FpML), http://www.oasis-open.org/cover/fpml.html

[28] DARPA Agent Mark-up Language (DAML), http://www.oasis-open.org/cover/daml.html

[29] XML-Based Ontology Exchange Language (XOL), http://www.oasis-open.org/cover/xol.html

[30] Extensible Financial Reporting Mark-up Language (XFRML), http://www.oasis-open.org/cover/xfrml.html

[31] Trading Partner Agreement Markup Language (tpaML), http://www.oasis-open.org/cover/tpa.html

[32] Artificial Intelligence Mark-up Language (ALICE), http://www.oasis-open.org/cover/aiml-ALICE.html