EXPERT SYSTEMS IN FINANCE – A CROSS-SECTION OF THE FIELD

Ljubica Nedović[1] and Vladan Devedžić[2]

[1]Čačanska banka ad Čačak,
Pivarska 1, 32000 Čačak, Yugoslavia
E-mail: ljubican@eunet.yu
Tel: +381-32- 302193, Fax: +381-11-225048

[2]FON – School of Business Administration, University of Belgrade
POB 52, Jove Ilića 154, 11000 Belgrade, Yugoslavia
E-mail: devedzic@galeb.etf.bg.ac.yu
Tel: +381-11-3950856, Fax: +381-11-461221

Abstract - This paper surveys a number of well-known expert systems in the domain of finance. The idea is to illustrate the span of typical expert systems in finance and to provide an insight into the approaches and techniques they employ. The systems considered in the paper are all working systems, and come from different fields of finance. Specifically, the systems considered are FINEVA (from the field of financial analysis), PORT-MAN (banking management), INVEX (investment advisory) and FAME (financial marketing). The paper also briefly presents DEVEX, an expert system for currency exchange advising in international business transactions. Financial institutions in less-developed and undeveloped countries that deal with currency exchange often face some specific problems that don't exist in well-developed regions of the world. DEVEX helps bank employees to cope with such problems. Since business transactions between developed countries and Third World ones represent an important part of the world's financial affairs, the problems covered in DEVEX go beyond local financial institutions.

Keywords - Expert systems, financial analysis, banking management, investment advisory, financial marketing, currency exchange advising, international business transactions.

1. INTRODUCTION

Financial experts possess knowledge acquired in practice and which cannot be found in literature or acquired in any other way, but which is invaluable to a business success of a firm or a financial institution. Therefore, special attention is paid to creating expert systems in a domain of finance, so that the domain knowledge could be accessible to a wider circle of people in the first place, and also to make the work in this field easier.

Important advantages in using an expert system are the uniformity of knowledge and a possibility of its improvement over time. For example, if an expert system is used in giving help while evaluating risks of investment in a firm, then every relevant parameter is treated with special attention, without the fear that some situation could be differently estimated by various experts or clerks, or that some important parameters would not be taken into consideration. If any new parameter is important for a company, then it is added to the knowledge base and is taken into consideration during the decision-making process.

There are several groups of expert systems for finance according to the problem they treat:

- *Expert systems for financial analysis of firms* (Smith & McDuffie 1996; Matsatsinis, Doumpos & Zopounidis, 1997; O'Leary, 1995). An advantage of expert systems is that besides quantitative ratios such as profitability, total profit, long term and short term dept and other, they can treat qualitative ratios also: position at the market, organisation of personnel, business reputation, marketing flexibility, etc. Successful financial analysis determines the firm's status, that is, a level of risks for a possible investment or a credit approval.

- *Expert systems for analysing the causes of successful or unsuccessful business development* (Apte et al., 1989; Pinson, 1992; Ruparel & Srinivasan, 1992; Chinn & Madey, 1997). Because of the possibility to draw a conclusion about a certain business development or a specific project by looking backward or forward, and also the ability to follow data that can change over time, it is possible to find causes of the temporary state of a fairs in the business/project and predict its future.

- *Expert systems for market analysis* (Chan, Dillon & Saw 1989; Dhananjayan, Raman, & Sarukesi 1989; Smith, McDuffie & Flory 1991). If a product has already been produced before by a company, then its sale can be analysed by an expert system. The system can take into account different factors that can possibly decrease the sale (high price, low quality, bad commercial, stiff competition, etc.), and then on the bases of the analysis the company should decide on further steps (to improve quality, to improve production, or to start producing a new product).

- *Expert systems for acquiring knowledge in a subfield of finance* (Boer & Livnat 1990; Brown & Wensley 1995; Hartvigsen 1992). These expert systems are widely used in educating managers and other financial experts. Besides, this kind of knowledge can advance and be improved rapidly, so the knowledge bases of such expert systems are suitable means for its conservation, improvement and reusability.

2. FINEVA - AN EXPERT SYSTEM FOR FINANCIAL ANALYSIS OF FIRMS

The complete methodology for knowledge acquisition and representation in the field of financial analysis is implemented in the system called FINEVA (FINancial EVAluation) (Matsatsinis, Doumpos & Zopounidis 1997). The FINEVA system is a multicriteria knowledge-based decision support system for the assessment of corporate performance and viability. The system has been developed using the M4 expert system shell, by N.F. Matsatsinis, M. Doumpos and C. Zopounidis of Technical University on Crete.

Financial analysis of firms involves identification of the strengths and weaknesses of firms, mainly through judgemental procedures concerning the qualitative evaluation and interpretation of financial ratios. The technology of expert systems (ES-s) technology is well suited to these kinds of tasks. The symbolic reasoning of ES-s enables them not only to draw conclusions, through a process similar to the one used by human experts, but also to provide explanations concerning their estimations.

In the development of FINEVA, the knowledge from the international literature has been used and further knowledge acquisition has been conducted through a series of interviews with the financial experts from a bank in Greece. Decision tables have been used to elicit the knowledge from the experts in the most efficient way, while a decision tree provides a graphical representation of the acquired knowledge. The representation through production rules is used to implement the acquired heuristics in a knowledge base.

The output that FINEVA produces is a specific ranking of the firms considered, according to a class of risk.
The assessment of corporate performance and viability is achieved through the combination of the evaluation of financial status and the qualitative evaluation of the firm, Figure 1. For the two basic components the experts proposed equal weights.

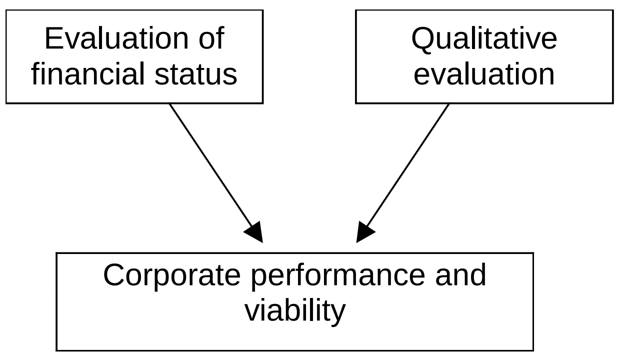

Figure 1. The assessment of corporate performance and viability in FINEVA

The financial ratios were categorised in three major groups - profitability ratios, solvency ratios, and managerial performance ratios. The profitability is evaluated through the examination of the profitability of total assets (industrial profitability) and stockholder's equity (financial profitability), the gross profit to total assets ratio, and the profit margin, Figure 2. The evaluation of solvency is accomplished by the analysis of the debt capacity (the short-term, the long-term and the global debt capacity) and the liquidity of the firm (the direct and the general liquidity), Figure 3.

The expenses (interest expenses, general and administrative expenses) and the mean period between accounts receivable and accounts payable are combined to evaluate the managerial performance of a firm, Figure 4.

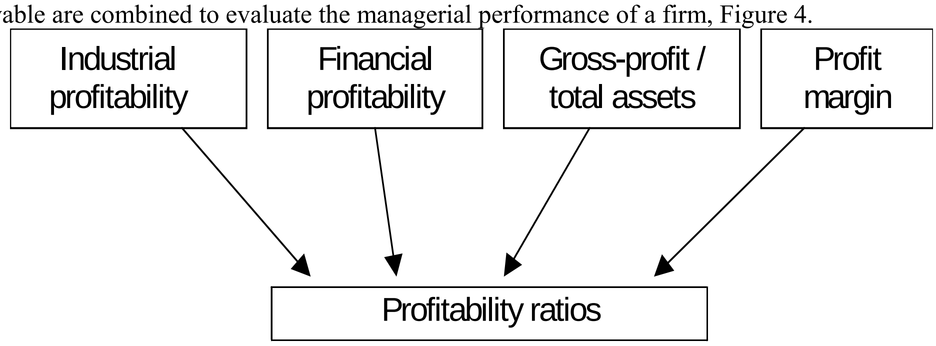

Figure 2. The profitability ratios

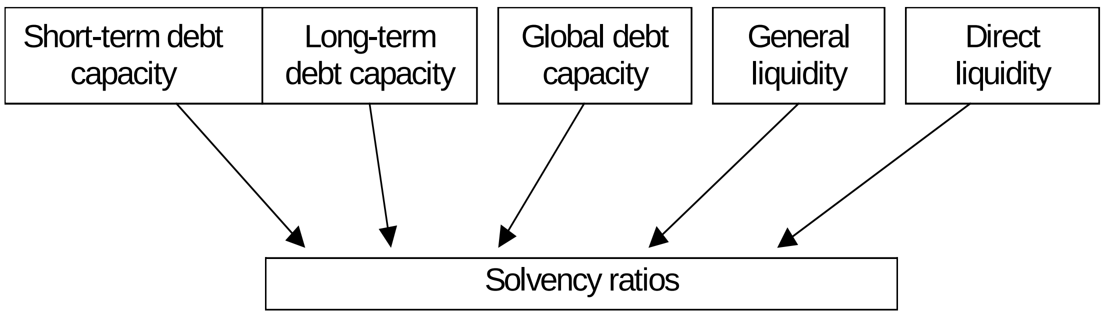

Figure 3. The solvency ratios

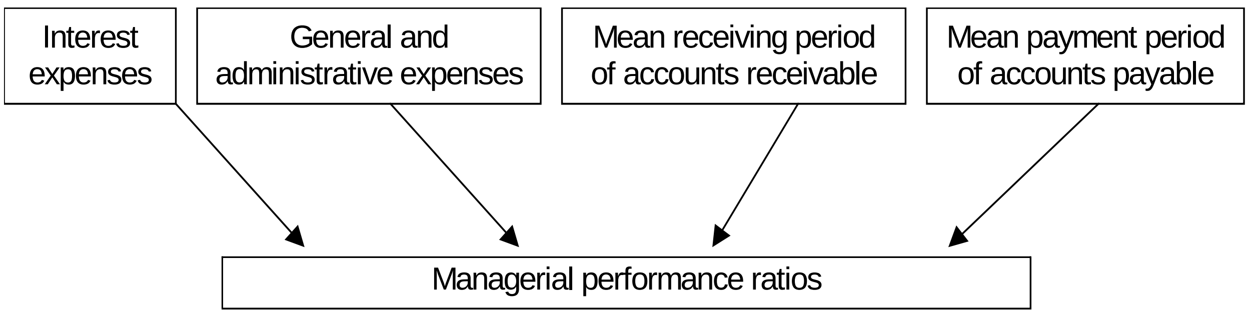

Figure 4. The managerial performance ratios

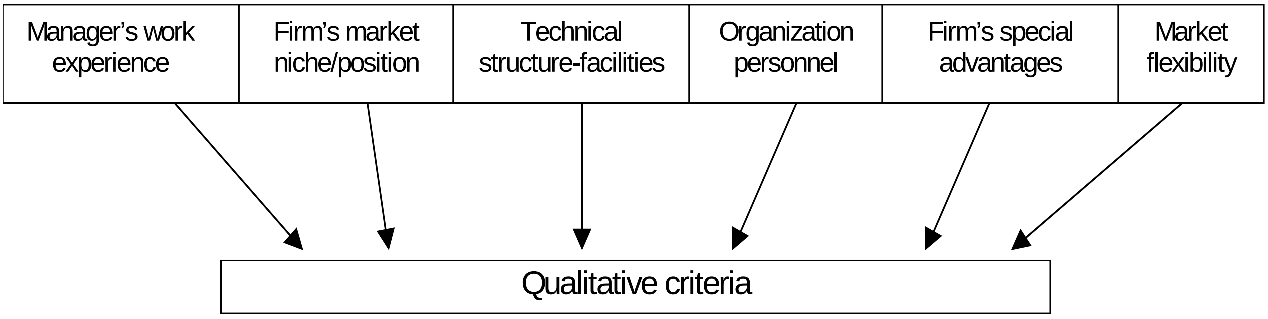

Figure 5. The qualitative criteria in FINEVA

The qualitative criteria for qualitative evaluation of the firm are: manager's work experience, firm's market niche/position, technical structure-facilities, organization personnel, firm's special competitive advantages, and market flexibility, Figure 5.

Each one of these criteria has been modelled using a five-point scale: not satisfactory, medium, satisfactory, very satisfactory and perfect. Figures 6 and 7 show some examples of how the above financial ratios and the qualitative criteria are represented in FINEVA.

Industrial profitability A1: A1 < 10% — not satisfactory 10% < A1 <= 20% — medium 20% < A1 <= 30% — satisfactory A1 > 30% — very satisfactory	**Financial profitability** A2: A2 <= 17.5% — not satisfactory 17.5% < A2 <= 20% — medium 20% < A2 <= 23% — satisfactory 23% < A2 — very satisfactory
Gross profit/Total assets A3: A3 <= 0% — not satisfactory 0% < A3 <= 50% — medium 50% < A3 <= 75% — satisfactory A3 > 75% — very satisfactory	**Profit margin** A4: A4 <= 0% — not satisfactory 0% < A4 <= 50% — medium 50% < A4 <= 100% — satisfactory A4 > 100% — very satisfactory
Short-term debt capacity B1: B1 < 25% — not satisfactory 25% < B1<= 50% — medium 50% < B1 <= 75% — satisfactory 75% < B1 <= 100% — very satisfactory	**Global debt capacity** B2: B2 > 80 % — not satisfactory 60% < B2 <= 80% — medium 40% < B2 <= 60% — satisfactory B2 <= 40% — very satisfactory
Long-term debt capacity B3: B3 <= 0.5 — satisfactory B3 > 0.5 — not satisfactory	**General liquidity** B4: B4 >= 2 — satisfactory B4 < 2 — not satisfactory
Direct liquidity B5: B5 <= 1 — not satisfactory 1 < B5 < 1.5 — satisfactory B5 >= 1.5 — very satisfactory	**Financial expenses** C1: C1 > 5% — not satisfactory 3% < C1 <= 5% — medium 2% < C1 <= 3% — satisfactory C1 <= 2% — very satisfactory

General and administrative expenses C2: C2 > 8% not satisfactory 6% < C2 <= 8% medium 4% < C2 <= 6% satisfactory 2% < C2 <= 4% very satisfactory C2 <= 2% perfect	**Received period of accounts receivable** C3 > C4 not satisfactory C3 <= C4 satisfactory
Circulation of inventories C5: C5 increasing not satisfactory C5 reducing or stable satisfactory	**Circulation of customers and notes receivable** C6 <= C7 satisfactory C6 > C7 not satisfactory

Figure 6. Modelling of the Financial Ratios in FINEVA

Managers' work experience:	
Negative experience	Not satisfactory
No experience	Medium
Positive experience up to 5 years	Satisfactory
Positive experience 5-10 years	Very satisfactory
Positive experience more than 10 years	Perfect
Firms' market niche/position:	
Strong competition, firm's weak position	Not satisfactory
Strong competition, established and competitive firm	Medium
Moderate competition, firm's strong position	Satisfactory
Weak competition, firm's leadership position	Very satisfactory
Single position, monopoly	Perfect
Technical structure-facilities:	
Old and inappropriate equipment, outdated production methods	Not satisfactory
Moderate technical structure, non-competitive production cost	Medium
Relatively modernized equipment	Satisfactory
Sound technical structure, full modernization scheme under way	Very satisfactory
Excellent structure, modern production methods	Perfect
Organization personnel:	
Lack of organization/staff hiring policy	Not satisfactory
Moderate organization/staff hiring policy	Medium
Moderate organization/staff hiring policy, willingness to improve	Satisfactory
Good organization/staff hiring policy	Very satisfactory
Excellent organization/staff hiring policy	Perfect
Firm's special competitive advantages:	
The firm does not possess expertise for its production methods	Not satisfactory
The firm possesses a small amount of expertise for its production methods	Medium
The firm possesses a satisfactory level of expertise for its production methods	Satisfactory
The firm possesses an exclusive expertise for its production methods	Very satisfactory
Market flexibility:	
The firm does not follow market trends, produces low-demand products	Not satisfactory
The firm has a limited flexibility	Medium
The firm has a satisfactory flexibility	Satisfactory
The firm follows market trends	Very satisfactory
The firm is a leader in its production branch activity	Perfect

Figure 7. Modelling of the Qualitative Criteria in FINEVA

The acquired knowledge is represented in the knowledge base through production rules. FINEVA's knowledge is represented in the set of total **1693** rules, representing over **13.000** possible combinations of the evaluation criteria. Basically, there are two major sets of rules: one concerning the evaluation of firms based on their financial ratios, while the other involves the qualitative evaluation of a firm through the examination of the strategic variables. The first of these sets is divided into three subsets of production rules including a set of profitability rules, a set of solvency rules (including liquidity rules), and a set of managerial performance rules. Each of these subset is further divided into smaller and more specific subset of production rules, Figure 8**.**

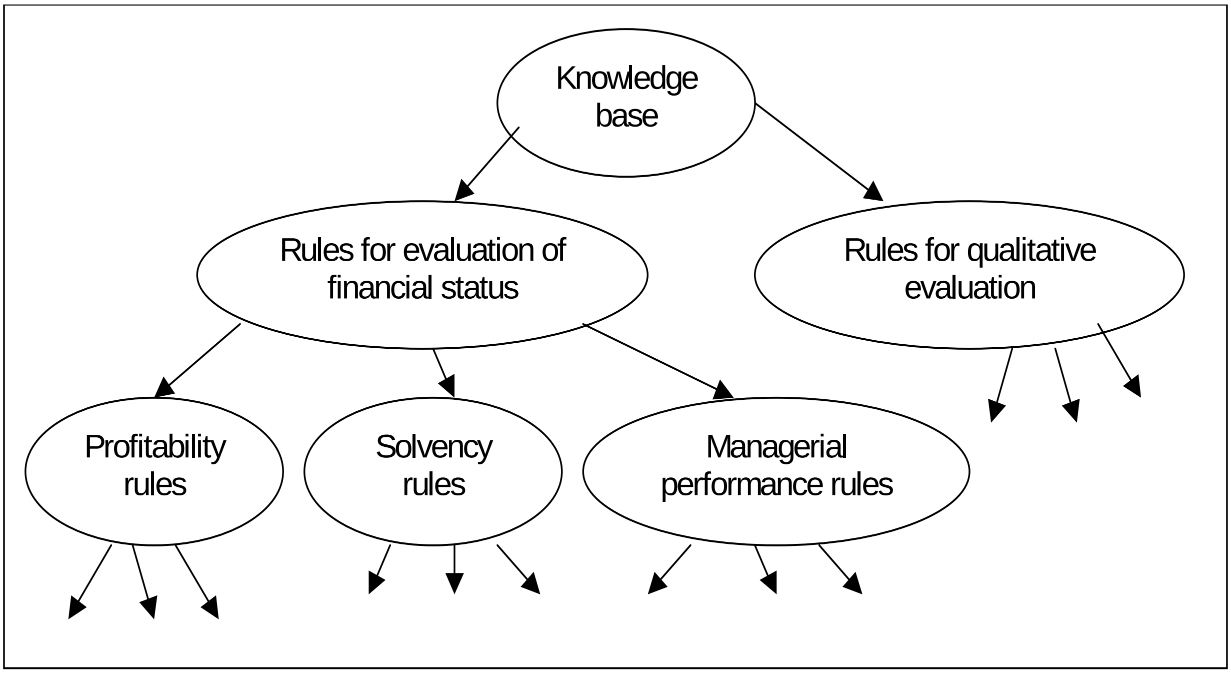

Figure 8.Knowledge base in FINEVA

Two concrete examples of FINEVA's rules are shown in Figure 9:

IF Direct-liquidity = satisfactory **AND** General-liquidity = not-satisfactory **THEN** liquidity = medium	**IF** Financial-status = very-satisfactory **AND** Qualitative-evaluation = satisfactory **OR** Qualitative-evaluation = very-satisfactory **OR** Qualitative-evaluation = perfect **THEN** Expert-system-evaluation = very-satisfactory

Figure 9: Some rules from FINEVA's knowledge base

Along with the production rules, some meta-rules were also necessary to represent the heuristics that the experts use in several practical cases when assessing the performance of a firm, as well as to reduce the time needed to reach the final estimation. Meta-rules do not provide an estimation needed to draw a conclusion, but control and guide the inference process to specific sets of production rules or even the modification of production rules.

For example, in cases where the managers' work experience is 'not satisfactory' (negative work experience) then the expert assumes that the organization of the firm cannot be perfect. This meta-rule initially checks the set of possible (legal) value of the criterion organization, excluding the value 'perfect' (italics represent commands of the M4 expert system shell), (Figure 10):

IF *kbentry (L:legalvals*(organization)=[not-satisfactory, medium, satisfactory, very-satisfactory, perfect, unknown])
AND work-exp = not-satisfactory
AND *do(add L:legalvals*(organization)=[not- satisfactory, medium, satisfactory, very- satisfactory, unknown])
THEN *set legalvals*(organization).

Figure 10: An example of a meta-rule in FINEVA

In cases when the value of a criterion is unknown, the limited information is ignored by the system by directing the inference process to the corresponding set of rules which do not examine the unknown information. The accuracy of the conclusion reached depends on the amount of the available data:

Whencached(financial-expenses **is** unknown **or** general-and-administrative-expenses **is** unknown) = rule-1630.

FINEVA's inference engine that draws conclusions about the performance of the examined companies employs both the forward- and the backward-chaining method. The forward-chaining method is used to guide the inference process to a set of rules (meta-reasoning), and the backward-chaining method is applied within this set to derive a conclusion. This inference strategy closely reflects human-expert logic and decision making in the domain.

3. PORT-MAN – AN EXPERT SYSTEM FOR PORTFOLIO MANAGEMENT IN BANKS

One of the major areas of service provided by the banking industry is to help people to plan the financial aspect of their lives. In order to function effectively, banks must be able to advise their customers on the best possible arrangement that would suit their individual investment needs. This implies that the investment advisor must have knowledge of the products offered by the bank and the ability to recognize the customer's needs and match these needs with the appropriate products.

Currently this service is performed by the bank officers. One of the problems is the non-consistency of the advice given from these officers. Certain products could be well-known to some officers but are ignored by others. Hence, the same investment situation could lead to different advice from different bank officers.

Furthermore, the consultation process is usually complicated and may take some time. This delay could lead to the investor's impatience with the process, with the likely loss of the client to the bank. Hence an expert system can greatly improve a bank's service to its customer, as it could make this service more readily available and greatly speed up the process.

Port-Man expert system has been developed by Y. Y. Chan, T. S. Dillon and E. G. Saw at the La Trobe University in Bundoora, Australia (Chan, Dillon. & Saw 1989). Port-Man is a banking advisory system designed to assist bank officers to give advice on personal investment in a bank. It helps to speed up the consultation process and standardize the experience of the bank's financial consultants. The task of the system is to select a range of bank products that will satisfy the criteria for investment. The selected products are ranked according to the rates of return-on-investment and risk levels. Moreover, various side-effects for the investor, such as tax variation or pension adjustment, will be taken into consideration. Upon request, the system will give an explanation of how a product is selected. In addition, the user may query the system during the consultation process. Finally, Port-Man allows the user to change any previous input or investment criteria, and the system will then restart the process at the appropriate stage.

In general, the consultation process of Port-Man can be divided into 4 stages:

1) *Information acquisition*. Initially Port-Man acquires personal information about the investor. In case the investor has made a previous consultation with the system, Port-Man searches for the personal record of the investor from the data base. The subsequent consultation will then update the record accordingly. Otherwise, a new record will be created for the consultation. A minimum set of questions about the investment objectives and criteria will be asked via a form on the screen. More specific questions will be asked only if needed in later stages.

2) *Product selection*. In this stage, Port-Man searches for feasible products for the investment. Products are divided into different groups according to the product features, such as interest type, capital growth, etc. The product groups form a tree, with the most general group as the root node and the more specific sub-product groups as the successor nodes.

The search algorithm may then be considered as a tree search algorithm. Starting from the root node, the algorithm traverses to the successor nodes in a depth-first search manner. When a successor node is visited, the algorithm attempts to match the investment criteria with the features of the product group. If the matching succeeds, the algorithm continues its search at this node and at the successor nodes. If the matching fails, the algorithm prunes the branch from that node and backtracks to another unvisited branch of the most recent predecessor node. If the matching is undecided, the algorithm will request more information from the system, which will then collect the desired information by rules, if-needed functions, or direct queries to the user. Consequently, most interactive query-answer dialogs will take place at this stage.

No matter whether the matching is successful, unsuccessful or undecided, the system records the decision made and the corresponding justifications. Moreover, each piece of information will have a record of its own justification as well.

3) *Choice refinement.* Once the feasible products are found, the system considers the various side-effects that the products may have for the investors. Warning messages will be given if the investor's tax situation is affected or if his/her pension card entitlement is jeopardized by any of the selected products. It may even recommend the period of investment for some of the selected products. Finally, the system ranks the selected products according to the rates of return-on-investment and the risk levels.

4) *Explanation*. In this optional stage the user may question the system on how or why a product is selected. Recall that the system has recorded every decision and justification as it traverses along the product tree. A reverse list of the decision record may answer the 'how' question and the appropriate justification records could answer the 'why' question. In addition, the justification of the adjustment actions made in the choice refinement stage are also recorded so that more specific questions may be answered at this stage.

In Port-Man, frames are the major components of knowledge representation, while production rules are used to represent the control knowledge of product selection. System parameters, personal details of investors, investment criteria, and features of products are all represented in frames. To facilitate the system solution and to reduce the search space, the products with similar features are grouped together. Even the rules are grouped together and are attached to the appropriate frames. Rules are used to guide the system selection of the investment products and are attached to various slots in the frames. Hance, the control becomes modular and local to the frames. In general, Port-Man has six classes of frames:

1. Customer frames
2. Target frames
3. Product frames
4. Variable frames
5. Control frames
6. Objective frames

Customer and Target Frames: The personal record of each investor is internally represented by a customer frame. The slots consist of personal facts about the investor and a history of previous consultation with Port-Man. New information collected from the present consultation will also be inserted into the frame. For each consultation, a separate target frame will be created for the investor.

A target frame records a particular investment objective and investment criteria. The required product features and their justifications are all records in the target frame. In addition, a target frame has a set of attached rules to guide its search along the product tree.

Product frames: The products are classified into different groups, which form a hierarchical tree. Internally, Port-Man represents each product/product group by a product frame. A product frame is therefore a node in the hierarchical tree. It describes the common features of the product group. It also has a set of attached rules to guide the search for the sub-product groups. The general structure of a product frame is as follows:

Product group
'prod-trigger': *the chosen sub-product groups*
'rule-refine': *the set of refinement rules, useful for searching for the sub-product groups under this product group*
'sub-product': *list of sub-product groups referenced by this product group*
feature-1: *common features shared by the referenced specific products*
…
Feature-n: ...

Variable frames: Each system parameter has a corresponding variable frame. A variable frame is used to control and record how its value is derived. The rules required, the if-needed functions, the value derived and the context of the parameter are all stored and maintained in the variable frame. For instance, the parameter **Invest-Amt** has the following structure:

Invest-Amt:
'to-get-value' – 'query' <value facet>
- 'In order to determine the minimum amount the products should have' <justification facet>
'parent' - **Objective-2** <value facet>
'value' - unknown <value facet>
- **'Find out'** <if-needed facet>

The 'value' slot is used to store the value of the parameter. The 'Find-Out' is the name of an user-defined function used to derive the value of the parameter. The 'parent' slot indicates the dynamic relationship between the variable frame and another frame, and is particularly useful for explaining chains of "why" queries from the users. The 'to-get-value' slot is used to specify how the unknown value is deduced. The justification facet of the slot gives an explanation on why the variable is required by the system.

Control and Objective Frames: The control frame contains a list of tasks to be performed. A task is an action to be taken by the system and consists of a function that is to be executed. The name of the function is held in the control slot of the frame.

The objective frame is created dynamically by the system whenever Port-Man pursues a goal or performs a task. It has slots defining the purpose or the reason for the goal along with the name of either the variable frame or the previous objective frame from which the goal has originated. Hence in this way all the active objective frames are linked together. The objective frame will be deleted once its goal or task is completed.

The frame representation provides the system with a very modular control methodology. Most frames have some attached rules or functions. These rules or functions define the entire control within the attached frames. They also control the next frames to be activated. Hence the process can be considered as a sequence of frames activations. The search algorithm in the product selection stage is implicitly defined by the attached rules and functions in the product frames.

Maintenance comprises removing unwanted product frames, inserting new product frames, and updating existing product frames in the knowledge base. The search algorithm can also be specifically tuned for some particular product groups. A major problem encountered in large expert systems is the maintenance of the attached rules and the rules which activate the frames. For example, if a product is to be replaced by a new product, we have to remove the rule attached to the old product frame and allocate a new set of rules to the new product frame. Besides, we may need to update the context of the rules that activate the removed frame. A lot of searching will be involved, especially if the rule base is large.

Port-Man is implemented in XL, an expert system development environment developed by ISR, an Australian company. XL provides both syntactic and functional level manipulations of frames and rules. The tools provided by XL simplify the building and manipulation of frames and rules. Consequently, the knowledge engineers can focus on the implementation of domain knowledge. Frames in XL have the following structure:

Frame-name :
Slot-1 <value facet>
<certification facet>
<justification facet>
<if-needed facet>
...
Slot-n - ...
...

where Frame-name must be a symbol while Slot-i and facets may be symbols or strings. The value facets of a slot can be retrieved by a binary function **of**. Function

'parent' **of** Invest-Amt

returns the symbol **Objective-2**. However, if the value facet contains the logical value unknown, the function specified by the if-needed facet, if it exists, will be called automatically (XL has 3 logical values: yes, no, unknown). Therefore, the statement

'value' **of** Invest-Amt

results in calling the function **Find-Out,** whose value will be returned. Moreover, if the if-needed facet has missed a system event, the relevant event function (RELEVENT) will be called instead. The RELEVENT function may be redefined by the users. Therefore, it is useful to define a general Find-Out algorithm as the RELEVENT function.

XL allows all facets of a slot to be assigned by a XL statement called **remember**:

remember frame-name, slot, value [, certification ,justification ,if-needed].

The various facets of the 'value' slot can be easily assigned by the statement:

remember Invest-Amt, 'value',unknown,,,'Find-Out'.

Port-Man associates rules with the frame which defines the context in which the rules are applied. For example, if the rules are used to deduce the interest type, they will be grouped and attached to the **Interest** variable frame.

The grouping of rules is enhanced by other XL facilities. XL internally indexes rules by means of two master control frames (MCF-s) called the *lhs-frame* and *rhs-frame*. An element of an expression comprising the antecedent (condition) of a rule is stored as a slot name of the lhs-frame, while the rule name is placed in the value facet of the slot. Similarly, elements of the expression comprising the consequent (conclusion) of a rule are stored as slot names of the rhs-frame with the rule name in the value facets. Lhs-frame and rhs-frame have the following structures:

lhs-frame:
Slot-1 <value facet>
Slot-2 <value facet>
...
Slot-n - <value facet>

rhs-frame:
Slot-1 <value facet>
Slot-2 <value facet>
...
Slot-n - <value facet>

where *Slot-i* is an element of an expression, and *<value facet>* is a rule name. An element of an expression may be a constant, a string, a frame or a function name.

The frame properties of MCFs allow the rules to be grouped together by means of one expand statement:

expand [all] [lhs/rhs] relation [control-frame], pattern, object-name [,funct-name]

The keyword **relation** specifies that the expansion is with respect to frames. In brief, the **expand** passes the **object-frame** and all the information of the slot that matches in depth-first search manner the **pattern** of the **control-frame** (or **object-frame**, if control-frame is missing) to the user-defined function called **funct-name**. The keyword **all** ensures the expansion continues until all the slots are examined or that the **function** returns a logical value **no** at the top level. Hence, we may associate the rules whose consequent parts are relevant to an element, say 'fixed', to the frame called **Fix-Set** by the following statement:

expand all relation rhs-frame, 'fixed', Fix-Set, rule-filter

Thus, if a rule is given certain **keys** in a certain context, the grouping of rules can be made directly from the MCFs using the keys as the **pattern**. The maintenance can be also automated using the **keys**.

The explanation stage allows the user to examine the current chain of reasoning during the consultation and how the system has arrived at the solution. This stage is designed not only to justify the chosen products, but also to help the bank consultants to feedback their experience and to make recommendations to the system. In addition, Port-Man allows the user to change the values of the system parameters. In this way, it provides an efficient means to run the consultation using a different set of inputs.

4. INVEX - AN EXPERT SYSTEM IN THE FIELD OF INVESTMENT MANAGEMENT

Capital investment - deciding which product or business to support - is a very important business issue, because it is largely irreversible, usually involves long-term decisions, and affects the nature and the structure of the business.

The INVEX expert system (Vranes et al., 1996) helps the project analyst and investment decision-maker to determine whether a project is acceptable and, if it is, whether it is the best alternative, and to calculate the extent of the decision sensitivity to certain critical assumptions.

During a consultation, INVEX first asks about a customer's preferences and intentions, then builds up a customer profile, where the information asked from customers depends heavily on their intentions and the course of the consultation. These preferences and intentions are translated using production rules into the weights assigned to the different objectives in the multicriteria analysis knowledge source. INVEX is fed with data through the interface that most of the users already know - the spreadsheet. MS Excel plays the roles of a user-friendly, well-known and well-accepted front-end for data entry, a standard financial table generator and translator, and a client of the intelligent server based on the BEST tool for building expert systems (Vranes, 1992) performing background intelligent decision-making activities.

When all the input data (prepared in Excel sheets) are ready, INVEX performs the following steps. First, it divides the investments into five groups according to the values of dynamic parameters, Figure 11. Investments from the group VERY GOOD are accepted for the multicriteria decision-making (MCDM), while investments from the group VERY BAD are rejected. For the investments from the group GOOD, MEDIUM and BAD, a group-specific sensitivity analysis (GSSA) is performed and then a user is asked whether to accept or reject each of these investments. The first step decreases the number of investments that will take part in MCDM, by rejecting the bad choices. If specified, a risk analysis is performed on the accepted investments, and then the MCDM gives the optimal combination of investments for the given resources. Total preorder used in MCDM gives the best investments from the set of accepted investments. If the first step is skipped, then in the cases when we have only BAD and VERY BAD investments, the optimal combination would include the investments that may not be acceptable for the investor although they are the best in this case. Note that the optimal combination of investments need not include all the best investments from the total preorder.

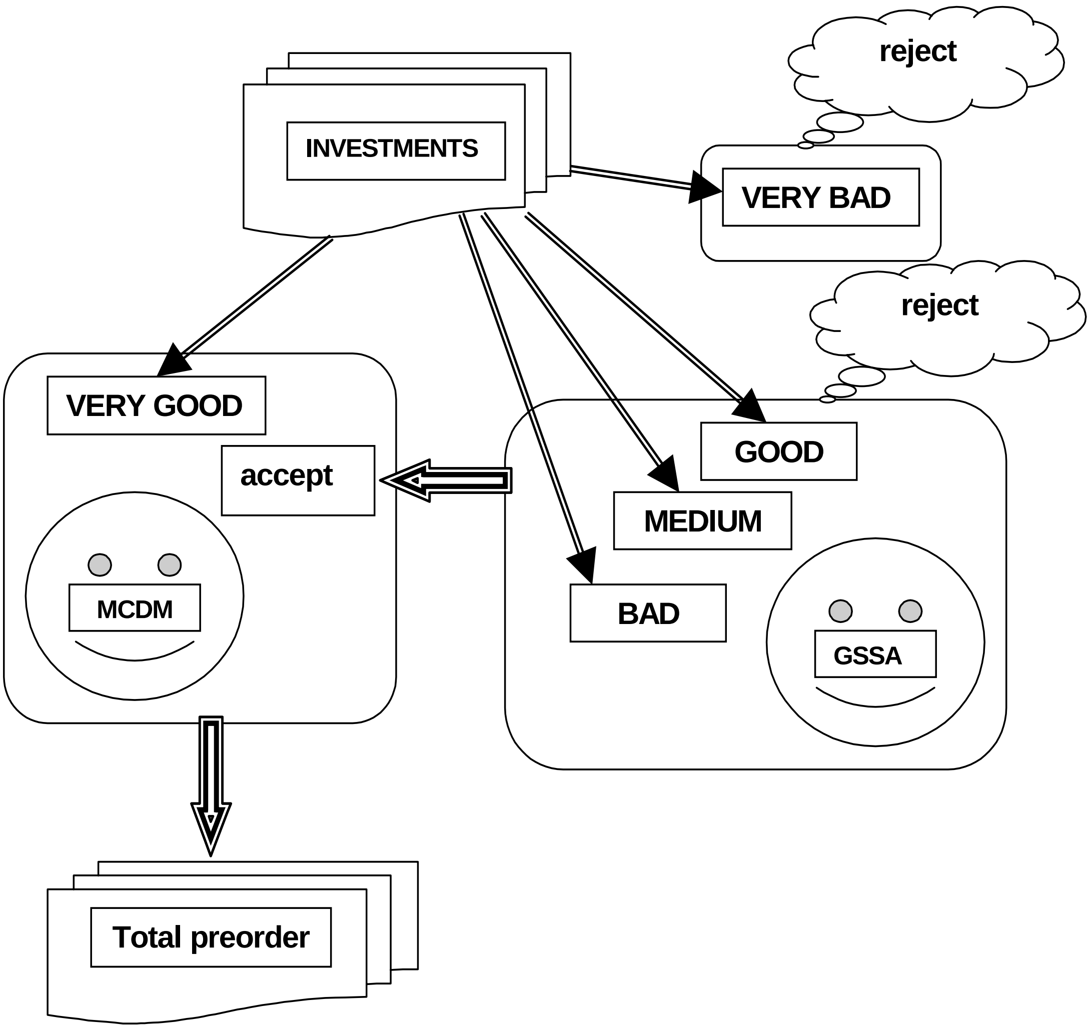

Figure 11: Division of investments in INVEX (after (Vranes, Stanojevic, Stevanovic & Lucin, 1996))

INVEX uses four dynamic indicators of project desirability:

- Relative net present value of investment (*v*);
- Return on investment (*r*);
- Payback period (*p*);
- Period of achieving the critical breakeven point in the exploitation of investment when the net flow becomes positive (*c*).

Financial acceptability of the project can be tested, and alternative project can be divided into five groups, depending on the extent of their desirability (Figure 12).

A project belongs to group **A** if:

- the net present value of investment is greater than the present value of invested resources ($v>1$);
- the return on investment is greater than 25% (assuming that the compound interest is 10%);
- the payback period p is shorter than the reference period $p0$ ($p<p0$);

- the period of achieving the critical breakeven point c is shorter than the reference period $c0$ ($c<c0$).

Such an investment is described as VERY GOOD.

An investment belongs to group **B** if:

- the net present value of investment is greater than 0 ($v>0$);
- the return on investment is between 10% and 25%;
- the payback period is shorter than the reference period;
- the period of achieving the critical breakeven point is shorter than the reference period.

An investment that belongs to group B can be described as GOOD.

Group **C** contains the investment whose:

- net present value of investment is small than 0 ($v<0$);
- return on investment is between 6% and 10%;
- payback period p is slightly longer than the reference period $p0$ ($p>p0$);
- period of achieving the breakeven is very short;

or:

- net present value of investment is greater than 0 ($v>0$);
- return on investment is between 10% and 15%
- payback period p is considerably longer than the reference period $p0$;
- period of achieving the breakeven is slightly longer than the reference period.

Group C contains investments that can be described as MEDIUM.

If an investment belongs to group **D** then its parameters have the following values:

- the net present value of investment is less than 0 and its absolute value comparable with the present value of invested resources ($v<-1$);
- the return on investment is less than 7%;
- the payback period is considerably longer than the reference period;
- the period of achieving the critical point is longer than the reference period.

An investment from D group is described as BAD.

Investments belonging to group **E**, have the following properties:

- net present value of investment is negative and its absolute value is greater than the present value of invested resources ($v<-1$);
- return on investment is less than 4%;
- payback period is much longer than the active life of investment;
- the breakeven point is never achieved.

An investment from group E is described as VERY BAD.

GROUPS	***v***	***r***	***p***	***c***
A	> 1	> 25%	< *p0*	< *c0*
B	> 0	> 10% and < 25%	< *p0*	< *c0*
C1	< 0	> 6% and < 10%	> *p0*	< *c0*
C2	> 0	> 10% and < 25%	>> *p0*	> *c0*
D	< -1	< 7%	>> *p0*	> *c0*
E	< -1	< 4%	>> *p0*	>> *c0*

Figure 12. The groups projects in INVEX and value its parameters.
Legend: ***v***- *r*elative net present value of investment,
r - return on investment,
p - payback period,
c - period of achieving the critical breakeven point,
p0 - reference payback period,
c0- reference period of achieving the critical point.

INVEX assumes that the user will definitely accept the investments from the group VERY GOOD and reject the investments from the group VERY BAD. For the investments from group GOOD, MEDIUM and BAD, where the user not sure whether to accept or reject a particular investment, the system performs a group-specific sensitivity analysis which gives some additional information that can help the users to make up their minds. Note that if the total preorder is used, it is not possible to give any additional information to the users that can help them to decide whether to accept or reject a particular investment.

Another perspective on investment decision-making relates to the issue of future uncertainty and its consequences for planning and decision-making. However, high returns are often associated with high risks. A major role of INVEX is to aid managers in assessing various future alternatives and the levels of risk and return associations between them. A complete knowledge source is dedicated to the risk-bearing attitude.

The decision-making process forces planners and analysts to assign values to uncertain future consumption and present investment in terms of present consumption. The process of risk analysis is employed to determine uncertainty in planning future investments and using all resources. The major sources of uncertainty are the price and demand projections, suggesting that major efforts have to be directed towards meaningful forecasting for cost and revenue variables.

In the risk-bearing model built into INVEX, a probability distribution is defined for investment effectiveness given a combination of uncertain variables that affect profits. In order to provide, for instance, net present value probability distribution, subjective probability distributions are first evaluated for the series of broad categories of revenue variables, cost variables, project life, the cost of capital and so on.

Figure 13 illustrates one of the charts produced by INVEX. It represents **BALANCE** w.r.t. net income.

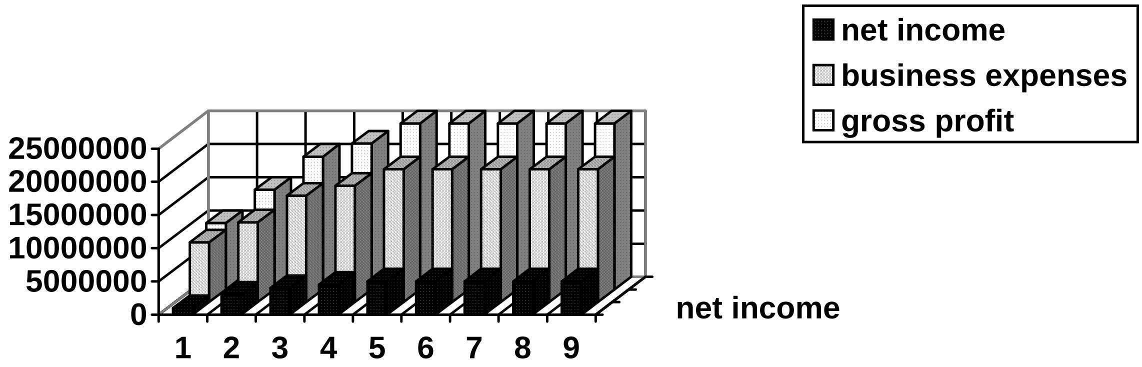

Figure **13. BALANCE** w.r.t. net income chart, as produced by **INVEX**

5. FAME – AN EXPERT SYSTEM FOR FINANCIAL MARKETING

Financial marketing represents a new and challenging domain for expert systems. Financial marketing is the activity that determines the most beneficial offer to a customer within an agreed set of financial parameters. Needless to say, the item must also satisfactorily address the customer's technical requirements and be competitive with other similar offers in the marketplace (active marketing situations).

Financial marketing is knowledge-intensive in nature. That is, not only does successful financial marketing require good marketing and financial skills, it also requires skills in mapping the technology being marketed onto the customer's requirements, and the ability to combine all these into meaningful, efficient actions, utilizing a vast amount of market data on product and services, historical trends, competition, and the customer's corporate financial profile. Given the high volume of information, problem solving has to deal very frequently with incomplete or uncertain scenarios. This naturally gives rise to multiple solutions based on varying assumptions.

This characteristic of financial marketing makes it necessary for any marketing proposal to be supported by extensive arguments for it to be of sellable value. There are usually no well-defined criteria for determining the best solution for a customer's problem. This domain is highly characterized by this lack of single answers, even from a single seller's viewpoint. It therefore becomes very important to be able to strengthen one's proposal by providing appropriate justifications and alternatives.

The FAME system is an expert system for financial marketing that has been developed long ago (Apte et al., 1989), but it still represents a good example of financial expert systems. Moreover, some ideas and techniques similar to those employed in FAME can be seen in modern intelligent systems as well (see the Discussion section).

FAME is a knowledge-based advisory system that helps in the preparation of comprehensive financial marketing recommendations for the mainframe computer business. It runs on Lisp workstations, and extensively utilizes the advanced I/O features that are commonly available on these workstations. The system operates as an interactive assistant, i.e., the user remains in complete control during a problem solving session.

FAME may be viewed as a large, heterogeneous knowledge-based system. The term heterogeneous refers to the nature of the knowledge used by FAME to solve typical problems in the domain. A variety of techniques are used to capture this knowledge, ranging from rule-based classification models and heuristic search algorithms to conventional and hybrid analytical techniques. To service this cross-section of problem solvers, so that they may communicate, share, and be controllable in an autonomous fashion, it is important that a model of the problem domain and the evolving problem solving situation be uniformly accessible to all the components.

Figure 14 illustrates the multi-layered architecture of FAME, which the FAME authors' experience has indicated to be very suitable one for large and heterogeneous knowledge based system. Thus, the FAME knowledge bases capture the essence of objects and their inter-relations as encountered in the domain, and relevant to intelligent problem-solving in it. For financial marketing in the mainframe computer business, one needs to model entities such as the products in the market today, their historical trends in terms of price and performance, the vendors who manufacture these products, the third parties who finance their use and acquisition, and customers. This broad base of knowledge about the domain is then usable by a variety of expert problem solvers. These expert sub-systems help a user in the interactive construction of a customer's financial profile, and subsequently, in designing marketing proposals. Among them, they offer a variety of automated services that can be utilized for generating competitive proposals and explanations, information walk-throughs, and related tasks.

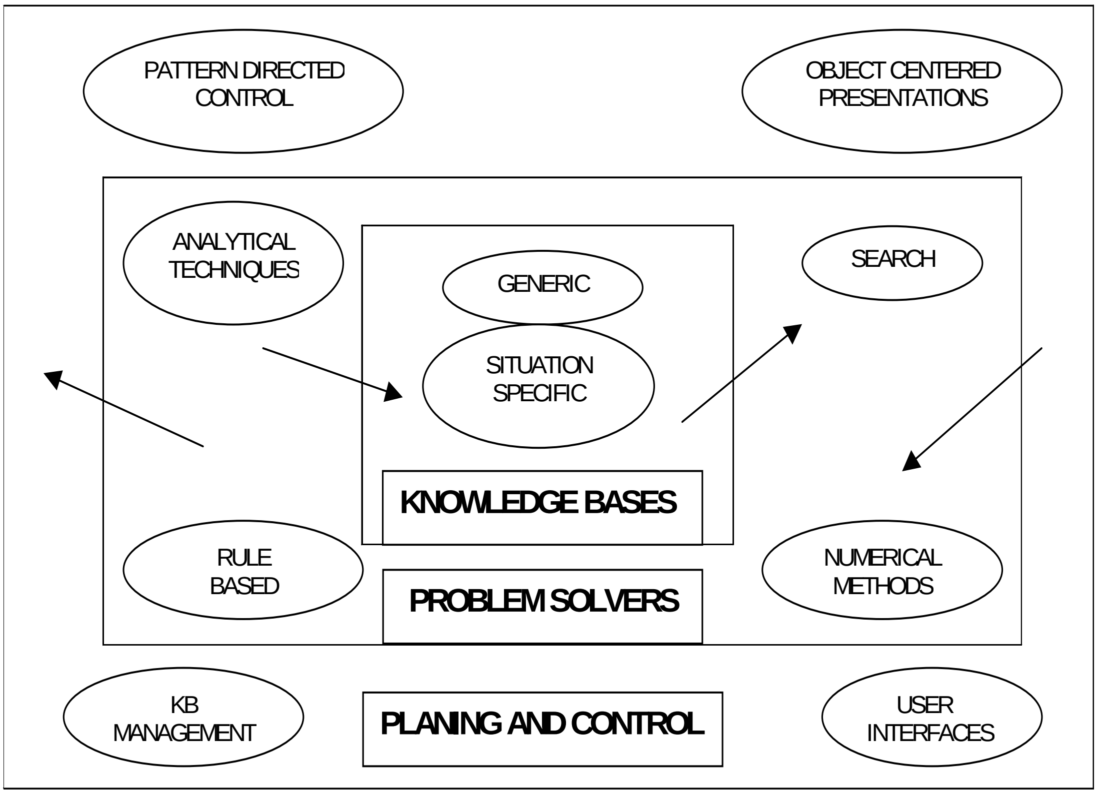

Figure 14. Multi-layered organization of the FAME system.

FAME used the K-Rep shell for knowledge representation. K-Rep views knowledge as a collection of objects in a structured inheritance network. K-Rep provides a mechanism for representing the very common and most natural characteristics of objects and their inter-connections using the fundamental algebraic relations of subsumption and attribution.

The primary object in K-Rep is called a concept. Concepts may be specializations of other concepts, in which case the more specific concept inherits attributes from the more general. Information about the concept's attributes is given via a binary relation (called *role relation*). K-Rep also provides a facility for performing definitional classification on new concepts. Using the classifier one can perform a restricted pattern matching by creating a new concept corresponding to the pattern, classifying the new concept, and retrieving those concepts which are its specializations.

This object-centered modelling allowed efficient access to tremendous amounts of inter-related knowledge that is typically required for producing detailed arguments and justifications for problem solving steps. Classifying the domain in terms of objects allowed the authors of FAME to quickly build hierarchies and taxonomies of objects by their different classes. These classes included, but were not limited to, financing mechanisms, manufactures, products (past, present, and future), customers, financiers, etc. as relevant to the market of mainframe computers. This categorization not only makes it easy to build such knowledge using structured inheritance networks, it also allows acquisition of such knowledge from experts either with a knowledge engineer and/or via computer-based acquisition tools.

Knowledge acquisition and maintenance are important issues for knowledge-based systems, even more so when we deal with very large knowledge bases. Thus the user may traverse or browse through the knowledge base by searching through the object's subsumption and attribution links. This link-driven traversing is extremely powerful for browsing large pieces of non-linear organizations, which is exactly what large knowledge bases are in their structure.

The FAME system consists of well over 2000 complex objects connected in a K-rep network of multiple inheritance links. This model permitted to build intelligent planning, control, and problem solving in FAME.

6. DEVEX - AN EXPERT SYSTEM FOR CURRENCY EXCHANGE ADVISING

The DEVEX expert system for foreign currency exchange advising in international business transactions is developed at Čačanska bank in Čačak, Yugoslavia. According to the classification described in section 1, DEVEX could be put into the group of expert systems for acquiring knowledge in a subfield of finance, but it can also be applied to concrete foreign currency transactions on the international level.

6.1. The model of foreign currency exchange in business transactions

Definitely one of the most important banking services and information-exchange products is SWIFT - the international system for electronic payments and sending messages. Financial institutions from all countries use SWIFT for currency exchange and other financial transactions. All banks that want to exchange information or perform banking transactions quickly and accurately are included in the SWIFT Net. SWIFT represents the world's standard in banking (S.W.I.F.T., 1998.).

Performing currency exchange with foreign countries over SWIFT Net imposes certain standards for message composing. These requirements must be built in the information systems of those banks that deal with currency exchange.

The participants that take part in international business transactions and a flow of information between them can be represented as the dependency graph in Figure 15**.** Client is a native physical or legal subject, or a foreign citizen. Home bank is the coordinator of services in which the clients have their foreign currency accounts, where they maintain their foreign currency deposits and through which they perform various transactions. Home banks or any other financial institutions that are authorized for such tasks have their own foreign currency accounts at foreign banks, through which transactions are performed at a client's request. Foreign firms, which act as importers or exporters of goods and services, also have accounts in a foreign bank, and that way they can participate in currency exchange of the home bank (indirectly). There may also be a connecting bank in the transaction chain between the home and the foreign banks. Through this connecting bank, the entire currency exchange between home and foreign banks is performed, or only a part of this exchange can be performed through it if a bank is not well-known in the international banking business market. For example, a home bank can only collect foreign deposits for its clients through the connecting bank.

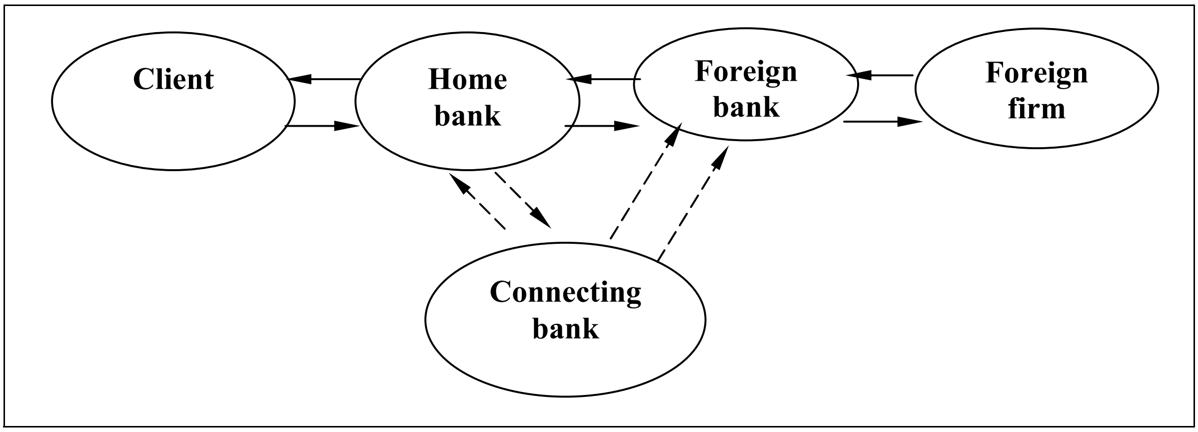

Figure 15. Participants in international business transactions

In some foreign currency business transactions all the participants mentioned above can appear, while only some of them participate in some other transactions. Payment to a foreign country (flow of resources) is a situation in which all participants take part. For example, if a client issues a demand for payment to a bank for a foreign firm, the home bank processes it and provides resources, and sends an order to a foreign bank to take a denoted sum of foreign currency off the accounts and pay it into the account of the foreign firm.

Buying currencies at the ICM (Inter-banking Currency Market) is performed only between the home bank and the connection bank: the connection bank processes all demands for buying foreign currencies issued by clients and sends them to ICM, usually through a mediator. For example, the mediator for Čačanska bank (in which DEVEX is installed) in such transactions is the National Bank of Yugoslavia (NBY). When a foreign currency is bought, NBY informs home bank which makes counter value payments (in local currency) and distributes foreign resources to the clients according to their demands.

6.2. The problem space of DEVEX

DEVEX is intended to help employees of home business banks who work on various tasks connected with foreign business transactions of a concrete bank. The main reason for building this system has been the fact that foreign business transactions represent a group of tasks that are very important in practice, but at the same time are extremely complex. It is very difficult to model them in a procedural, algorithmic way. The following discussion illustrates this statement, provides an insight into a part of complex heuristics associated with foreign business transactions, and explains the rationale for applying the expert system technology.

Foreign business transactions are a part (a subsystem) of the global information system of a bank or any other financial organisation in general. The subsystem is responsible for several tasks (transactions) that are performed within the country or with a foreign country, and which are very specific either in documentation or in the ways of performing them. Specifically, the tasks of this subsystem include currency exchange with foreign countries, various payments in foreign currencies for clients, buying of currencies at the Inter-banking Currency Market, and the like. Successful completion of such tasks affects business transactions of a firm itself and also the business and economy of the entire territory that this bank covers. Hence a special attention is paid to these tasks while developing financial information systems.

Besides, foreign exchange transactions deal with a foreign currency and its equivalent value in local currency on the day of exchange, which is later used for making various statements. An additional problem is the fact that the home bank can work with several different exchange rates (such as selling foreign currency for cash, buying foreign currency with cash, selling and buying foreign currency through bank accounts, the medium rate, and so on), but only one can be used in a specific transaction.

Different kinds of business transactions are involved and they could be performed in different ways. The main problem is to choose the most efficient way to perform a transaction. For example, in payments to a foreign country DEVEX considers several factors in detail:

- What sources ensure payments? Does the firm have its own deposit at the account? If not, then the firm can purchase foreign currency from the bank itself, from some other firm that runs its business through the same bank, or at the ICM, depending on its temporary equivalent value of what? or other business interests.

- What currency is the most profitable one? If a firm makes its payments in DM, but it has US$ and CHF on its account, then the firm will sell the currency that is the most profitable w.r.t. DM.

- Through which foreign bank the payment will be made? This effects the time needed for making the payment. The most convenient foreign bank is the one that is the nearest to the foreign firm to which the payment will be made, and at which the home firm has its account. It is also important whether the home bank possesses the preferred currency, or a conversion to the needed currency must be done. Another issue is whether the firm has the currency at some other foreign bank when the resources will be transferred, whether the foreign firm has an account at that bank so that the resources could be transferred, or another bank (or banks) in the chain where the foreign firm has an account must be looked for. In that last case, it is important that this bank has its account at the foreign bank through which the home firm is making the payments.

Besides the knowledge of accounting (from what account the payment can be made and in what situations, what accounts will be booked, what accounts have demand or debt balance status, etc), the knowledge of legal regulations is also needed in this field: whether a payment in cash is possible and when, on what basis the resources can be kept, whether the a firm is registered for a certain payment transaction, the exact amounts of commissions and other expenses. Furthermore, this is an area where legal regulations have constantly been changing and they must strictly be obeyed. For example, in Yugoslavia every currency transaction with a foreign country is under strict control of the Federal Board for Documentary Control.

Documentation that is needed for a certain transaction depends on the way of performing the transaction. If the bank possesses the preferred currency, only the payment order and the order for compensation are needed; otherwise, the home bank must have additional orders for conversion, transfer, etc.

If several orders for payments to foreign countries have arrived, and at the moment the bank has limited resources, insufficient to cover all orders, then it is the job of an employee to decide which orders will be covered and which ones will be delayed (to decide on priorities). The employee can make decisions according to different criteria: whether the usage of goods is limited, or shipment must not be delayed, or maybe the import of a repro-material effects the work of several firms or an entire industrial branch, or maybe a firm imports machines in order to improve production in the future (for several months or a year; in that case the payment can be delayed), etc.

The orders that get processed can contain errors. An employee has to estimate if it is just a trivial mistake that he can corrected personally, or he must arrange for a telephone consultation with an employee from the firm, or if he has to send the order back to the firm. It is possible that the document itself does not allow any corrections (a letter of credit, for example).

This knowledge can be applied only to currency transactions of a bank that also deals with other services, such as counter work, citizen's savings, dealing with economy, etc.

Given all the complexity of dealing with foreign currency exchange described above, it follows that it is very hard to make a procedural program that would make decisions how to perform some transaction, prepare orders, and book for or transfer resources from an account to another one. Procedural programs in this domain are suitable only for currency accounting or booking of (already prepared) orders and drawing out necessary statements on the basis of booked items.

6.3. Implementation

If several orders for payments to foreign countries have arrived, and at the moment the bank has limited resources, insufficient to cover all orders, then it is the job of an employee to decide which orders will be covered and which ones will be delayed (to decide on priorities). The employee can make decisions according to different criteria: whether the usage of goods is limited, or shipment must not be delayed, or maybe the import of a repro-material effects the work of several firms or an entire industrial branch, or maybe a firm imports machines in order to improve production in the future (for several months or a year; in that case the payment can be delayed), etc.

The decision criteria built in DEVEX have been contributed by financial experts. Figure 16 shows the final set of financial and non-financial criteria-parameters built in DEVEX, which are used as evaluation criteria for the assessment of the priorities for payments. The contribution of each criterion has been modelled using an integer value on the scale from 0 to 10, where 0 is the minimum and 10 is the maximum value. Concrete values depend on the accuracy of payments and other data about business firms in a previous period.

Name	Description	Code
Financial criteria		
Possesses resources on account	The firm has sufficiently large balance on its account in the foreign currency required for the payment	**A**
Bank can make payments	The business bank has sufficiently large balance on its account in a specific foreign bank, in the foreign currency required for the payment	**B**
The firm's importance for the bank		
	The amount of the firm's deposit (percentage of its capital) in the bank	**C1**
	The bank's income acquired from that firm in the previous period (the bank's commission for already realized financial transaction with foreign countries)	**C2**
	How regularly the firm pays he commission to the bank	**C3**
Non-financial criteria		
Type goods for which the payments are being made	Depends on the goods' persistency: short-persistency goods, medicament, longer-persistency goods, unlimited-persistency goods	**D**
Purpose of imported goods	Goods for mass-consumption (deficit goods), repro-material, machines and tools for production increase	**E**
Firm's importance for the region where it belongs		
	The firm's effect on the operation of other firms	**F1**
	The number of people from the region that work for the firm	**F2**

Figure 16. Criteria for determining the priorities of payments in DEVEX

The acquired knowledge is represented in DEVEX' knowledge base as a set of 320 production rules. The system has been developed using the EXSYS Professional tool.

The following example shows the production rule that calculates the value of the parameter C2 from Figure 16:

IF (average number of the firm's transactions <=5 **AND**
the bank's commission acquired from the firm's transactions >= 10.000) **OR**
(average number of the firm's transactions >=10 **AND**
the bank's commission acquired from the firm's transactions <= 10.000)
THEN C2:=3

DEVEX' inference engine draws conclusions concerning the priorities for payments by applying such rules from its knowledge base to the input data about specific payments.

7. DISCUSSION

Summarizing the previous sections, Table 1 shows most relevant details about the systems described. In spite of the fact that it shows only five different financial expert systems, it can give a good flavour of the diversity of domains, approaches, design, tools, and techniques used in this area.

Table 1. Summary of the five systems described

System	FINEVA	Port-Man	INVEX	FAME	DEVEX
Domain	Financial analysis	Portfolio management in banks	Investment management	Financial marketing	Currency exchange advising
Output	Ranking of firms, according to a class of risk	Select a range of bank products that will satisfy the criteria for investment	Determine whether a project is acceptable and, whether it is the best alternative	Recommend changes to a marketing proposal	Assessment of the priorities for payments
Used at	The ETEVA bank in Greece	The ASK bank in Australia	N/A	IBM US Marketing & Services	The Cacanska banka IN Yugoslavia
Tool/Shell	M4	XL	BEST	K-Rep	EXSYS Professional
Knowledge acquisition and representation	Decision tables, decision tree, rules, meta-rules	Production rules, frames, hierarchical tree	Fuzzy set, Excel spreadsheet, concept	Heterogeneous system, semantic networks, concepts, binary relation	Production rules
Reasoning	Forward chaining, backward chaining	Grouped and attached rules in the control and objective frames, depth first search manner	Multiple criteria decision-making, sensitivity analysis, demons associated on slots	Hybrid analytical techniques	Forward chaining, backward chaining
Uncertainty	N/A	N/A	Uncertain-random variable, rough model and fine-grain model	N/A	N/A
Size	**1693** rules	N/A	N/A	2000 objects	320 rules
Status	Usable system	Usable system	N/A	Usable system	Prototype
Reference	Matsatsinis et al., 1997	Chan, Dillon. & Saw 1989	Vranes et al., 1996	Apte et al., 1989	

Two other important issues of financial expert systems deserve special attention. The first one is adherence of design to current trends in the broad area of intelligent systems. Published information about financial expert systems indicates that most of them are stand-alone systems, developed using different shells. Most of such shells have relatively closed design - they have been developed only with typical ES philosophy in mind, without taking a broader context into account. By this broader context we mean general software architecture philosophy, increased modularity, object-orientation, component-based systems, ease of integration with other systems in distributed environments, embedded systems, knowledge sharing and reuse, and other currently popular and important issues of software engineering and knowledge engineering.

In order to partially illustrate these current trends and the broader context of intelligent systems, consider Figure 17. It shows key ideas of OBOA, a systematic approach to design and development of software for intelligent systems (Devedzic & Radovic, 1999). It is based on a multilevel, general object-oriented model of intelligent systems, and has been developed starting from some observations regarding difficulties in making extensions to methods and tools for building intelligent systems - scalability efforts - as well as in reusability and embedding-oriented efforts. The OBOA approach makes clear distinction between generic, low-level intelligent software components, and domain-dependent, high-level components of any intelligent system. Each software component of an intelligent system is defined as belonging to a certain level of abstraction (Figure 17-a): the *primitives* level (where components are, e.g., logical expressions, frame slots, rule clauses, and neurons), the *units* level (things like rules, frames, neural networks, fuzzy sets and their combinations), the *blocks* level (sensors, planners, controllers), the *system* level (e.g., financial expert systems, assembly lines, robotic cells, intelligent medical systems), and the *integration* level (multiple intelligent agents, distributed systems). Components at the lower three levels are domain-independent, while components at the upper two levels are domain-dependent. Also, components at each level can be specified along several dimensions, such as knowledge representation, reasoning, knowledge acquisition, etc., Figure 17-b. An intelligent system (e.g., an expert system) in a traditional vertical domain (e.g., finance) can be specified at the *system* level, or even at the *integration* level, as being based on some components at that level plus a number of components from the lower levels. If the components at each level are specified based on ontologies developed at that level - which is one of the current trends in, e.g., The Semantic Web - then the system's generality, scalability, reusability, and embedding capabilities greatly increase. Note that only one of the five financial expert systems described here - FAME - partially adheres to such a design. It is object-oriented, views knowledge as a collection of objects in a structured inheritance network, and hence it is rather natural and easy to extend and reuse.

Level of abstraction	Objective	Semantics
Level 1	Integration	Multiple agents or systems
Level 2	System	Single agent or system
Level 3	Blocks	System building blocks
Level 4	Units	Units of primitives
Level 5	Primitives	Parts of units

(a)

Level of abstraction	Dimensions D_1	D_2	. . .	D_n
Level 1				
Level 2				
Level 3				
Level 4				
Level 5				

(b)

Figure 17. Levels of abstraction and dimensions in the OBOA model

The other important issue is that of commercialisation. It is interesting to note that very few financial expert systems are offered as commercial products. Most often such systems are developed for a specific customer, or a financial institution develops the system in-house. It is highly unlikely that such a system will make its way as a commercial product. On the other hand, that can be considered a shame, simply because many financial expert systems are potentially useful to many other customers. For example, systems like DEVEX can be used with only minor changes in many banks in undeveloped countries, since the currency exchange problems faced in business transactions of such countries are much alike.

Some exceptions to this rule do exist, however. For example, the Stock Guru system, developed and used in Netherlands and United Kingdom, is offered on the Web at US$ 169 (see http://www.stockguru.nl/, or http://www.stock-guru.co.uk/). Stock Guru is an expert system which uses complex digital signal processing algorithms to generate clear buying and selling signals. It can display price charts of all kinds of stocks and indices, plotted along with the so called Guru Lines (see Figure 18). These Guru Lines are useful for predicting when a security has peaked or bottomed or is continuing in its current trend. Stock Guru can be used as either a standalone system or a "second opinion decision maker" accompanied by daily performed technical analysis.

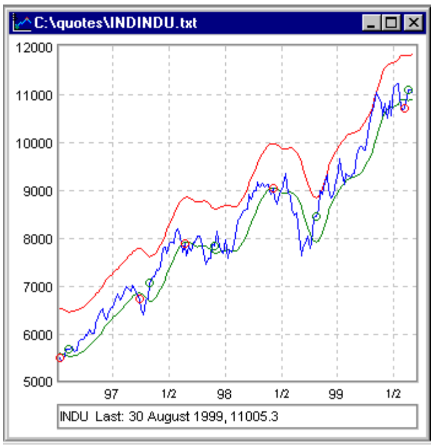

Figure 18. A plot from Stock Guru

8. CONCLUSIONS

The breadth of application domains of financial expert systems is best seen in surveys of the entire field of expert systems. One such survey from mid-1990s (Durkin, 1996), has shown that the number of financial expert systems developed and actively used in practice at that time was well over 300 and that the applications ranged from various banking subdomains (such as credit card application processing, evaluation of financial conditions of banks, security transaction analysis, and loan advisory) to bidding and bid preparation, financial planning, market analysis, tax advisory, portfolio management, allowance planning, stock market prediction, investment advisory, insurance risk assessment, and claims authorization and processing.

The five systems analysed in this paper have been picked for a typical cross-section of the field, focusing on the techniques, tools, approaches, problems, and practices used in development of financial expert systems. Rule-based technology still seems to dominate, but other techniques and approaches are present as well, especially connections to commercial database management systems.

In the forthcoming years, we expect financial expert systems to converge more with global trends in the broad field of intelligent systems (such as component-based architectures, agent-oriented applications, and embedding), as well as to turn to commercialisation and get more integrated into e-commerce and e-business, since many problems covered by specific financial expert systems are common in many financial institutions. A necessary prerequisite in that sense is the development of reference architectures and ontologies for various financial applications.

9. REFERENCES

Apte, C., Griesmer, J., Hong, S.J., Karnaugh, M., Kastner, J., Laker, M., & Mays, E. (1989). 'Utilizing knowledge intensive techniques in an automated consultant for financial marketing. *Expert Systems in Economics, Banking and Management*, (pp. 279- 288). Pau, L. F. (ed). Elsevier Science Publishers, B.V., North-Holland.

Boer, G.B., & Livnat, J. (1990). Using expert systems to teach complex accounting issues. *Issues in Asccounting Education***,** Spring, 108-119.

Brown, C.E., & Wensley, A. (Eds) (1995). Special issue on expert systems in accounting, auditing, and finance. *Expert Systems With Applications*, 9(4), 433-608.

Chan, Y.Y., Dillon, T.S. & Saw, E.G. (1989). Port-Man – An Expert System of Portfolio Management in Banks, *Expert Systems in Economics, Banking and Management*, (pp. 87-96). Pau, L. F. (ed). Elsevier Science Publishers, B.V., North-Holland.

Chinn, S.J., & Madey, G.R., (1997). A Framework for Developing and Evaluating Expert Systems for Temporal Business Applications. *Expert Systems with Applications*, 12(3), 393-404.

Devedzic, V., & Radovic, D., (1999). A Framework for Building Intelligent Manufacturing Systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C - Applications and Reviews*, 29(3), 402-419.

Dhananjayan, R.S., Raman, V.S.J., & Sarukesi, K. (1989). Application of expert system techniques for analysing firm's fall in market share. *Expert Systems in Economics, Banking and Management*, (pp. 289-308). Pau, L. F. (ed). Elsevier Science Publishers, B.V., North-Holland.

Durkin, J., (1996). Expert Systems: A View of the Field. *IEEE Expert* (pp. 56-63).

Hartvigsen, G. (1992). Limitations of knowledge-based systems for financial analysis in banking. *Expert Systems with Application*, 4, 19-32.

Matsatsinis, N.F., Doumpos, M., & Zopounidis, C., (1997). Knowledge Acquisition and Representation for Expert Systems in the Field of Financial Analysis. *Expert Systems with Applications*, 12(2), 247-262.

O'Leary, D.E., (1995). AI in accounting, finance and management. *International Journal of Intellligent Systems in Accounting, Finance and Management*, 4, 149-153.

Pinson, S., (1992). A multi-expert arhitecture for credit risk assessment: The CREDEX system. In D. E. O'Leary, & P. R. Watkins (Eds). *Expert systems in finance* (pp. 37-64). Oxford: Elsevier Science Publishers.

Ruparel, B., & Srinivasan, V., (1992). A dedicated shell for designing expert credit support systems. *Decision Support Systems*, 8, 343-359.

Smith, L.M., McDuffie, R.S. & Flory, S.M. (1991). A GAAP-based expert prototype for business combinations. *Financial and Accounting Systems*, Summer, (pp. 17-22).

Smith, L.M., & McDuffie, R.S., (1996). Using an Expert System to Teach Accounting for Business Combinations. *Expert Systems with Applications*, 10(2), 181-191.

S.W.I.F.T., (1998). User Handbook Standards, Brusseles, Belgium.

Vranes, S., Stanojevic, M., Stevanovic, V., & Lucin, M. (1996). 'INVEX: investment Advisory Expert System. *Expert Systems*, 13(2), 105-119.

Vranes, S., (1992). Expert System Shell Flexibility: BEST Case Study, in S. Tzafestas(Ed.) *Engineering Systems with Intelligence*, Kluwer Academic Publishers, (pp. 33-38).

Biographies (Vitae)

Ljubica Nedovic is a software engineer at "Cacanska banka, ad Cacak" of Cacak, Yugoslavia. She has graduated from Department of Computer Science, School of Mathematical Sciences, University of Belgrade, Yugoslavia, in 1987, and is just about to finish postgraduate studies at the School of Electrical Engineering, University of Belgrade, Yugoslavia. In the beginning of her engineering career, she used to do traditional SQL-based programming for database applications. Later on, she has extended her interests towards artificial intelligence, intelligent information systems, and expert systems. Her current interests include application of artificial intelligence techniques in the domain of finance, as well as object-oriented modelling and design, Java, Delphi, and the Internet.

Vladan Devedzic is an associate professor of computer science at the Department of Information Systems, FON - School of Business Administration, University of Belgrade, Yugoslavia. He has received all of his degrees from The School of Electrical Engineering, University of Belgrade, Yugoslavia (BS, 1982; MS, 1988; PhD, 1993). His main research interests include software engineering, intelligent systems, knowledge representation, ontologies, intelligent reasoning, and applications artificial intelligence techniques to education and medicine. So far, he has authored and co-authored about 160 research papers published in international and national journals and conferences. His major long-term professional goal is a continuous effort to bring close together the ideas from the broad fields of intelligent systems and software engineering. He has developed several practical intelligent systems and tools, and actively participates as a consultant to several ongoing projects in industry.