GENIUS™ Automated Underwriting System: Combining Knowledge Engineering and Machine Learning to Achieve Balanced Risk Assessment

David C. Golibersuch
GE Corporate Research and Development
PO Box 8
Schenectady, NY 12309
golibers@crd.ge.com

Rebecca Towne, Cheryl A. Wiebe
GE Capital Mortgage Corporation
6601 Six Forks Road
Raleigh, NC 27615
rebecca_towne@smtpmail.gecmc.ge.com

Abstract

The GENIUS Automated Underwriting System is an expert advisor that has been in successful nationwide production by GE Mortgage Insurance Corporation for two years to underwrite mortgage insurance. The knowledge base was developed using a unique hybrid approach combining the best of traditional knowledge engineering and a novel machine learning method called Example Based Evidential Reasoning (EBER). As one indicator of the efficacy of this approach, a complex system was completed in 11 months that achieved a 98% agreement rate with practicing underwriters for approve recommendations in the first month of operation. This performance and numerous additional business benefits have now been confirmed by two full years of nationwide production during which time some 800,000 applications have been underwritten. As a result of this outstanding success, the GENIUS system is serving as the basis for a major re-engineering of the underwriting process within the business. Also, a new version has recently been announced as an external product to bring the benefits of this technology to the mortgage industry at large. In addition, the concepts and methodology are being applied to other financial services applications such as commercial credit analysis and municipal bond credit enhancement. This paper documents the development process and operational results and concludes with a summary of critical success factors.

Introduction

Industry and Business Overview

In 1993, the US mortgage industry provided approximately $1 trillion for long term residential loans. This enormous industry is on the verge of major change as new technologies are being applied to improve responsiveness and efficiency (Schneider 1994).

Mortgage insurance is designed to protect lenders and subsequent investors (the secondary market) from unacceptable loss in the event of borrower default (Dennis 1992). Such insurance covers the gap between the amount of the property value the lenders and investors are willing to underwrite and the amount covered by the borrower's equity (initially the down payment). Consequently, it mitigates the higher risk of default associated with low borrower equity. This risk can change as real estate markets change. In an appreciating market, the probability of equity loss is minimal as the home can usually be sold for an amount sufficient to cover the gap. A depreciating market presents the opposite effect.

Mortgage insurance supports two important societal goals. First, it encourages a ready supply of capital to the home market. Second, it enables borrowers with limited equity to acquire mortgage financing. The major mortgage insurance suppliers in the US are: Amerin Guarantee Corporation, Commonwealth Mortgage Assurance Company (CMAC), GE Capital Mortgage Insurance Corporation (GEMICO), Mortgage Guaranty Insurance Corporation (MGIC), PMI Mortgage Insurance Company (PMI), Republic Mortgage Insurance Company (RMIC), Triad Guarantee Insurance Corporation, and United Guaranty Residential Insurance (UGI).

GE Capital Mortgage Insurance Corporation (GEMICO), a subsidiary of GE Capital Mortgage Corporation (GECMC) has been in the industry since 1983, and is an industry leader with a market share of more than 20%. During peak volume periods, GEMICO underwrites some 2,000 mortgage insurance applications per day, each representing approximately $25,000 in exposure.

The mortgage insurance industry is dynamic and highly competitive. Prompt and effective customer service is the most important competitive element. New products are introduced frequently often stressing the boundaries of acceptable risk. Price is an element of competition leading to constant pressure on production costs. In addition, recent legislation has defined certain practices as discriminatory. These regulations must be carefully considered to achieve fair lending and equal credit.

Underwriting and the Need for a New Approach

A central figure within this dynamic environment is the underwriter. The underwriter has the responsibility to determine if a particular application represents an acceptable risk given company policies, the borrowers' financial profile, and the property's current and anticipated future value. Traditional underwriting practices have evolved over many years with high reliance on individual experience. Adequate knowledge of many different loan types and markets is difficult to acquire. It has thus become an increasing challenge for underwriters to keep up with the numerous industry changes and continue to make prudent decisions while under constant pressure to reduce the time devoted to each application especially during periods of high volume.

In response to this situation, GEMICO management decided to reengineer the underwriting process and to support the new process by developing and deploying an expert adviser. The challenge was to develop a new approach that would support more responsive underwriting while continuing to properly account for all inherent risks. The new approach had to be able to respond quickly to changes in products and customer needs while ensuring consistency across all experience levels of individual underwriters and geographic regions.

Objectives of an Expert Adviser

High Quality Decisions The most important requirement was to promote high quality decisions that would lead to lower losses. This could be achieved by enabling underwriters of all levels of experience and in all geographic regions to emulate the best practices of experts to produce well-documented decisions consistent with business policy. While GE Capital Mortgage Insurance has published underwriting guidelines, a loan outside these guidelines can represent a prudent risk. Thus it was necessary for the recommendation to be based on an overall balanced assessment of risk factors as well as circumstances that mitigate that risk; this is in contrast to exclusive reliance on a rigid set of rules or guidelines. At the same time, it remained necessary to take account of the relatively few hard and fast guidelines that may not be excepted. It was important that recommendations be based on knowledge including both lessons from historical portfolio and delinquency experience as well as anticipated future considerations. Consistency across geographic areas was important as interstate banking laws continue to change. A loan in Texas may be submitted to an underwriter in Boston but the answer should be the same as if it were reviewed in Texas. Finally, decisions must be rendered in a manner consistent with all fair lending legislation.

Improved Customer Service This was to be achieved by facilitating faster responses to insurance applications as well as accurate and consistent communication of the basis for decisions.

Underwriter Productivity An important goal was to improve underwriter productivity. The amount of time could be reduced by quickly identifying those applications that did not require extensive review. For more complex or problematic applications, the system should guide the underwriter to those specific aspects of the application that are of concern. To support these goals the system must be a "glass box" in the sense that recommendations are clearly explained in terms meaningful to the underwriter. Additional productivity improvements were expected from reducing the time to train new underwriters as well as to update the knowledge of experienced underwriters.

Adaptability It was essential that the system be easily modified in response to new and changing products and changes in the industry, real estate markets, the economy, and business goals. Furthermore, easy modification was needed to allow timely incorporation of new information derived from ongoing analysis of the portfolio. To support this adaptability it was essential that the knowledge incorporated in the adviser was in a "glass box" in the sense that it could be easily accessed and understood by risk managers. As a consequence, they would be able to anticipate how the system would behave prior to implementation.

Efficacious Development Lastly, it was necessary to provide a rapid prototyping capability to ensure that the developing system was faithfully incorporating the business knowledge and policies and satisfactorily addressing the objectives discussed above.

Extensive use of the powerful concepts developed and made practicable over the past two decades by AI researchers was seen as the only feasible course to realize these many demanding objectives.

Prior Work

For many years GEMICO had been using a mainframe based system for processing and underwriting of mortgage insurance applications. This system provides an interface between several hundred nationally distributed personnel and underwriting support, reporting functions, and a central data base. In addition, this system has used a "Guidelines" module to identify characteristics of an application that signal risk concerns and so inform the underwriter. Also, it has long been the practice of GEMICO to track approved applications as they age and to perform statistical analysis of those that become delinquent or enter default to identify salient factors that predispose an application toward problems.

An early predecessor project was undertaken to replace the COBOL Guidelines program and its complex nested "if" structure with a rule based implementation to improve maintenance productivity. This re-engineering project was successfully completed but saw only limited production as the Guidelines approach was superseded by the overall risk assessment described in this paper (Still et al. 1991). It did, however, provide valuable experience in building and implementing a rule based system within the business's mainframe environment.

As early as 1988 risk management personnel had concluded that it would be preferred to replace the Guidelines module with an expert adviser that could provide an overall balanced risk based recommendation and eliminate or at least minimize specific guideline type exceptions. Early attempts were made to develop such a program using machine learning techniques such as neural nets. However, these efforts did not meet the business objectives for various technical and business reasons and were not pursued.

As a result of these prior activities, the GENIUS project had a substantial "head start". Business systems personnel were familiar with expert system tools. Risk factors were well understood. The project objectives were refined and considerable thought had been given to the structure of the underwriting decision process.

The Knowledge Base

Alternatives Considered

We will now briefly discuss alternative technical approaches considered to meet the challenging business goals. This will be followed by a more complete description of the reasoning architecture and the knowledge engineering process we employed.

Brute Force A straightforward approach would be to simply catalog all possible combinations of application attributes and associated recommendations for subsequent table lookup. Given the approximately one hundred attributes to be considered, each of which could take on multiple values, the number of possible combinations was estimated to be on the order of 10^{24} rendering such an approach totally impracticable.

Scoring A common technique widely practiced in the consumer credit card industry is statistical scoring where historical performance data for many borrowers are analyzed to determine the anticipated credit behavior of a particular applicant profile (Thomas, Crook, and Edelman 1992).

Pro

- Highly developed methodology.
- Extensive consumer credit industry experience.
- Provides a quantitative estimate of the probability of default.

Con

- Difficult to accommodate mitigating circumstances commonly encountered in complex applications such as a home purchase.
- Not conducive to explanation of recommendation beyond numerical score and therefore does not provide customer service advantages.
- Based on retrospective historical data; difficult to accommodate future oriented knowledge such as anticipated changes in the economy.

Traditional Knowledge Engineering In this earliest and still common method, heuristic expertise obtained from interviews with one or more experts is represented in the form of rules (Talebzadeh, Mandutianu, and Winner 1994). These rules reflect a common form of human reasoning; namely, when certain conditions are satisfied one makes corresponding conclusions and/or takes appropriate actions.

Pro

- Used to develop many successful systems.
- Good for capturing articulated knowledge such as policies and exceptions.

Con

- Difficult to capture the "compiled experience" of experts due to communication barriers and psychological problems with introspection.
- Large number of rules can be difficult to maintain.

Machine Learning Numerous machine learning methods such as induction and neural nets have been developed with the goal of resolving the knowledge engineering bottleneck (Weiss and Kulikowski 1991). They share a common starting point; namely, historical examples or sometimes examples of data and corresponding decisions prepared by experts. Depending on the particular method, an algorithm is then applied to more or less automatically infer and represent the knowledge implicit in the examples.

Pro

- Preparation of examples does not require the active intervention of a knowledge engineer.
- Good to capture "compiled experience" knowledge implicit in the examples.
- When appropriately applied can improve knowledge engineering productivity.

Con

- Not well suited to represent general knowledge and policies; as a result large numbers of examples may be required leading to another bottleneck.
- Not conducive to a good explanation facility.
- Maintenance may require extensive retraining, especially if the number of examples is large.

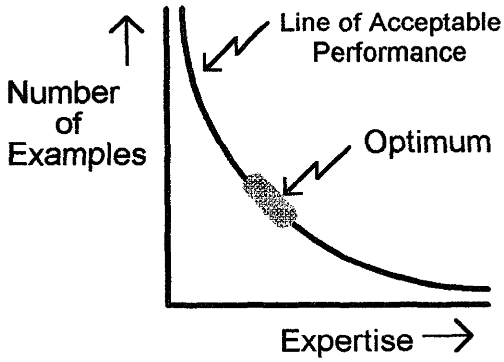

Figure 1. Hybrid Knowledge Engineering Tradeoff

Hybrid Knowledge Engineering This approach is designed to combine the advantages of traditional knowledge engineering and machine learning while minimizing their disadvantages. It is based on a judicious tradeoff between example based implicit knowledge and explicit expertise as depicted schematically in Figure 1. To obtain acceptable performance, a pure example based approach requires a very large number of examples that is often impracticable. Combining the example based approach with even a modest amount of explicit knowledge can substantially reduce the number of examples required and thus the overall development time. As one continues down the curve by adding more explicit knowledge fewer examples are required. However, as we continue to add more and more explicit knowledge, development time again becomes long and impracticable.

Evidential Reasoning Evidential Reasoning starts with a hierarchical description of the decision process wherein each node of the hierarchy represents an intermediate or final business consideration and opinion (Golibersuch 1995). Each node or sub problem contains a number of attributes describing the business factors considered for that node. Each attribute has a number of possible values. The attribute values are converted to numerical "evidence" values. The evidence for all attributes in a node is then combined and this combined evidence is used to determine the opinion for the node. Combination of evidence values is accomplished using a non-linear algorithm that has been found to emulate well the process of human "evidential" reasoning and is adapted from the work on certainty theory (Buchanan and Shortliffe 1985). This opinion is then propagated to the next higher level node where it becomes the value for the appropriate attribute in that higher level node.

The key to this approach is the appropriate choice of evidence for each attribute value and the interpretation of the various levels of combined evidence. In effect, the knowledge engineering process needs to answer the following questions. How much evidence should be assigned to an individual business fact? What is the relative importance of lower level nodes in reaching an opinion on a higher level node? How much combined evidence is required to reach a specific conclusion or opinion?

Pro

- Well suited to hierarchical business decision models.
- *As judged by experts*, provides an excellent emulation of the expert's ability to combine "apples and oranges", that is, disparate facts and intermediate opinions.
- Does not require numerous rules to explicitly describe how different combinations of facts and intermediate opinions are to be combined.
- Is fully compatible with traditional production rules. Thus rules can be used for those aspects of the reasoning process where they are the easiest and most natural method to represent the knowledge.
- Highly conducive to a complete and transparent representation of the knowledge as well as a flexible explanation facility, that is, "glass box".

Con

- Determination of the required evidence values can be a time consuming trial and error process.

Example Based Evidential Reasoning (EBER) For each node the expert provides a representative set of examples of attribute values and corresponding intermediate opinions. A semi-automated algorithm is then applied. The inputs to this algorithm are the examples. The outputs are a table associating a specific value of evidence for each attribute value and a table associating the combined evidence for the node to its corresponding opinion (Golibersuch 1995). Generally, a subset of possible examples, typically 5-10%, for each node is adequate.

Based on the substantial advantages, a hybrid approach incorporating knowledge engineering and Example Based Evidential Reasoning was chosen. Figure 2 is a schematic of this approach. During development, the method enhances knowledge engineering productivity by enabling the system to "learn" from example decisions of expert human underwriters and risk managers. In production use, analogous to a trial jury, the system weighs the "evidence" in a mortgage insurance application based on this "training" and is thus able to render a balanced, risk-based recommendation. Use of this approach during production and development will be further described in the following sections.

Reasoning Architecture

Decision Hierarchy A central feature of the GENIUS system is a hierarchical organization of the considerations used to determine a recommendation. The hierarchy used at the time of initial pilot production is shown in Figure 3. Each box represents a sub problem or node used to formulate intermediate and final opinions. Each node contains several attributes. For example, the Assets node includes attributes such as Borrower Equity, Number of Months Reserve, Source of Funds, and Non-borrower Contributions. Each attribute has multiple possible values. For example, Months Reserve is entered as one of four ranges: > 4, 2 to 4, -0.5 to 2, and < -0.5 months. Depending on the specific attribute values, an opinion will be deduced for each node. For example, the Assets node will have an opinion of Excellent, Good, Fair, or Poor. Opinions of lower level nodes serve a dual role as attribute values for higher level nodes. For example, a "Good" opinion for the Assets node will be propagated to the Applicant History node as the value of the Assets attribute for that node.

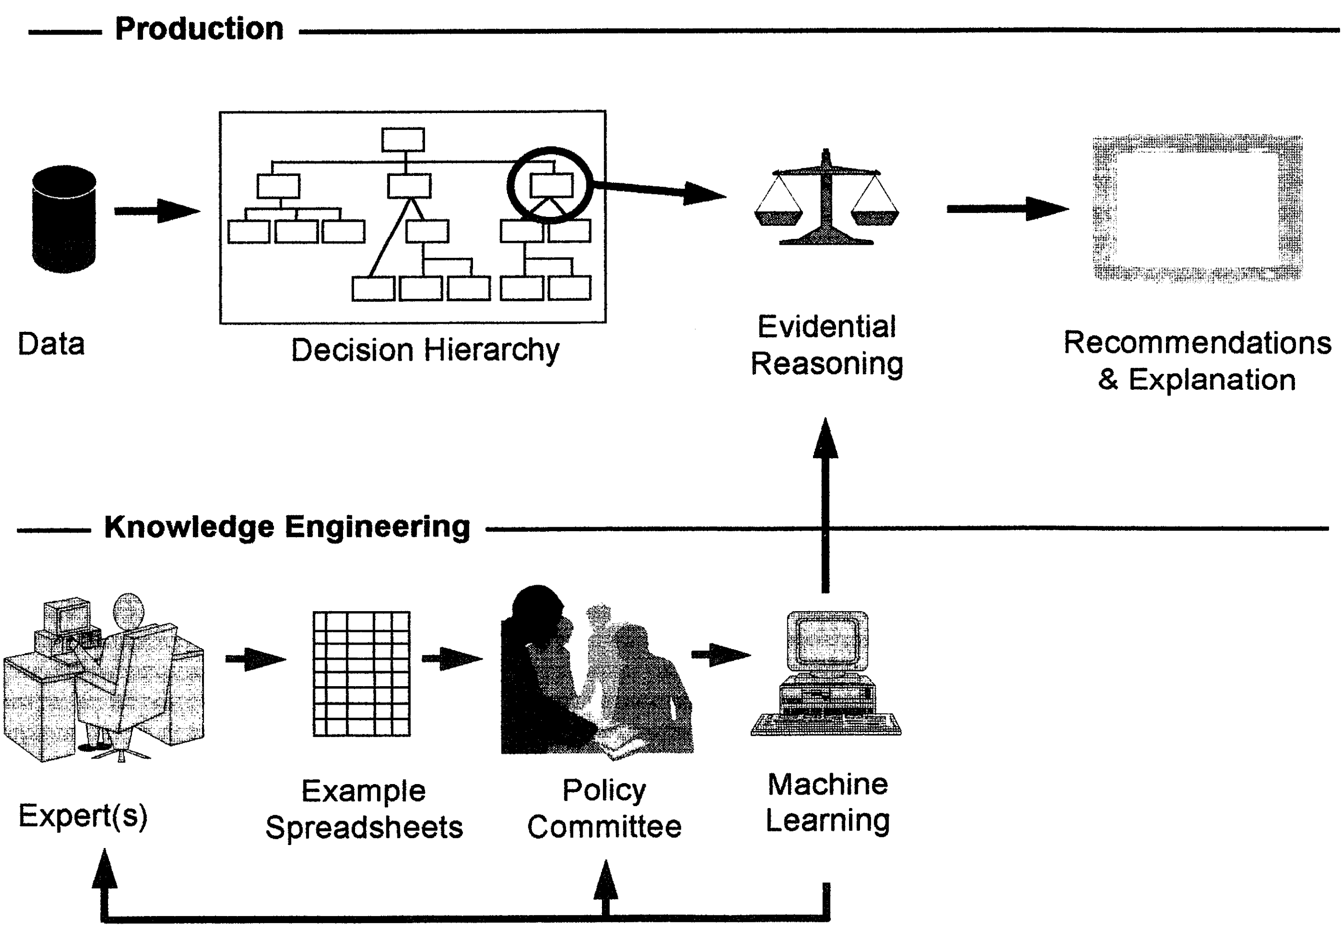

Figure 2. Example Based Evidential Reasoning

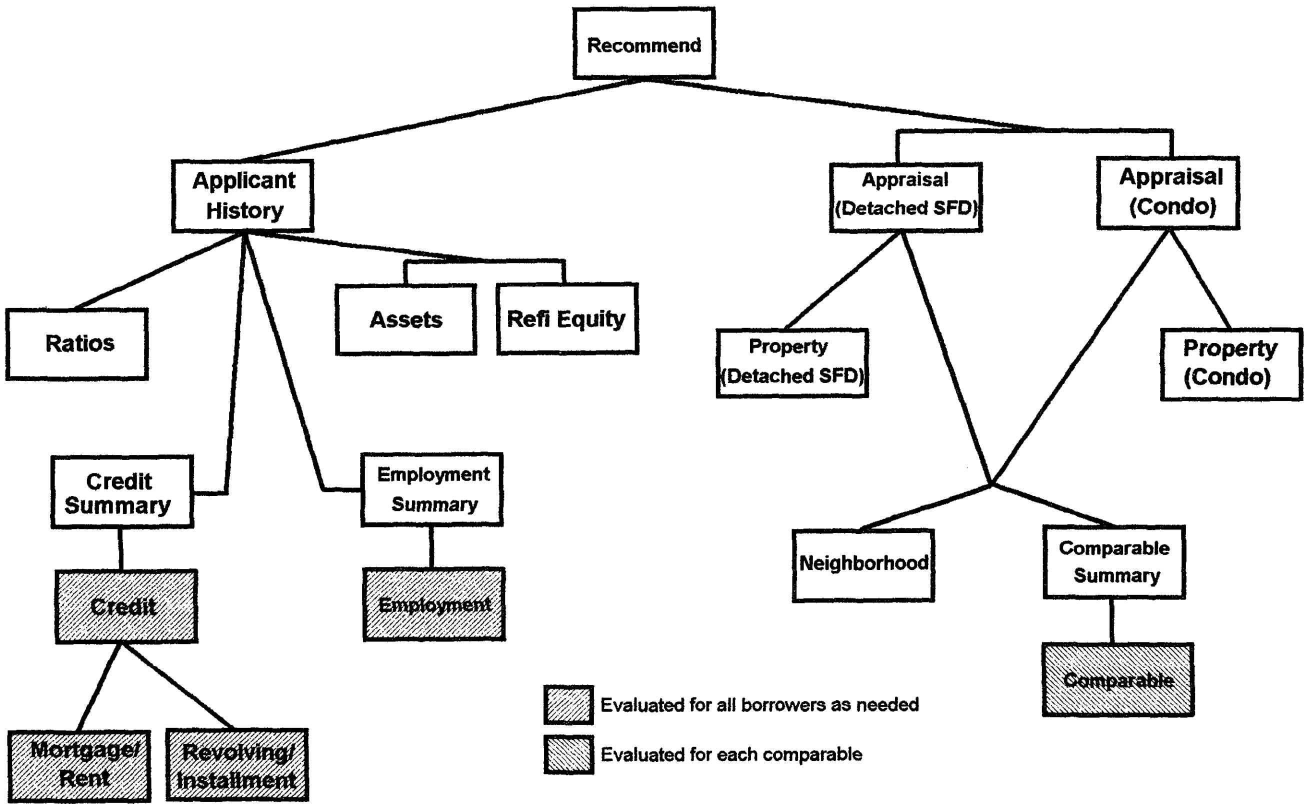

Figure 3. The GENIUS decision hierarchy.

The hierarchy is specifically configured for each application based on factors such as number of borrowers and application type, for example, purchase or refinance. After configuring the specific hierarchy to be used for an application the massaged data are transferred into the hierarchy from the client program. The hierarchy in GENIUS is implemented using an object oriented approach.

Evidential Reasoning Once a value has been entered into the appropriate attribute value slot, the evidence associated with that value is obtained by a table lookup. When the evidence is available for all attributes of a given node, the combined evidence is computed. The combined evidence is then used by another table look up to obtain its interpretation or opinion for the node. Subsequently this opinion is propagated to the next higher level node and the process is repeated.

Recommendation After all nodes have been evaluated the Recommend node will contain the overall risk based recommendation. The GENIUS system provides one of four recommendations: Approve, Minor Review, In-depth Review, and Decline.

Explanation In most commercial underwriting situations and in particular for mortgage insurance the underwriter's ultimate goal is to prepare a "considered opinion". This phrase expresses the underwriter's goal to provide an opinion supported by its basis and rationale. Some reasons why this is required include management of policy exceptions, identification of mitigating circumstances, and marketing considerations. It is especially important for the mortgage insurance industry where clear and consistent justifications can be a distinguishing competitive factor with customers. (This is to be contrasted with some business applications such as consumer credit assessment for bulk mailings where a simple opinion or "score" indicating likelihood of default but without explanation is adequate.)

In the GENIUS system the same screen that displays the overall recommendation contains a concise summary of the basis for that recommendation. The opinions for each of the major sub problems such as Assets and Comparable Summary are shown. For those cases where the node opinion is unfavorable, immediately below the node opinion are listed specific application concerns which have led to the unfavorable opinion. For example, if Assets is Poor the display might list comments such as

INADEQUATE EQUITY and INSUFF CASH/CLOSE. Most of this information is supplied by a generic explanation engine. Specific problem comments are stored in the same tables that are used for the evidence values. In this manner the explanations are easily accessed and maintained without the need for code revision.

Recall that a goal of the GENIUS system was to remove specific warning messages in favor of an overall risk based recommendation. This goal was largely achieved. However, there remain some situations that must be brought to the underwriter's attention. For example, a bankruptcy filed over three years ago will always be brought to the underwriter's attention even if the applicant's overall credit is favorable Such situations are handled by specific exception rules that effectively override the generic explanation engine. This design compromise was deemed acceptable as the number of such exceptions is small.

Benefits Of Approach

The technical approach chosen to implement the GENIUS knowledge base offers several important advantages:

- It leads to development of a robust and maintainable knowledge base that, as judged by experts and extensive field experience, works exceptionally well in a production environment.
- The "glass box" accessibility of the knowledge and basis for recommendations makes it credible and thus encourages acceptance among users and risk managers.
- It supports an optimal combination of machine learning to capture and represent implicit expert knowledge and a decision hierarchy and traditional rules to represent explicit knowledge and exceptions. This enables significant reductions in the time and cost of knowledge engineering as well as maintenance.
- It is applicable to a broad range of processes of interest to the financial services industry such as underwriting, credit analysis, product selection, and deal structuring. As a consequence, productive development of new applications is facilitated by design and code reuse.

Development and Deployment

Project Resources and Organization

The project team included representatives from three areas: GEMICO Risk Management, GEMICO Information Systems, and GE Corporate Research and Development.

The Risk Management team included a full time underwriting expert, and managers of portfolio risk, underwriting policy, collateral review, market analysis and risk policy. Management and coordination of the overall project were provided by this team, particularly the underwriting expert (author Wiebe) and portfolio risk manager (author Towne). In addition, a task force representing each region of the country to be served met at the beginning of the project and periodically throughout its course to provide guidance and review progress from the crucial user perspective. The total effort of the Risk Management team was approximately 2 person-years.

The Information Systems team was responsible for creating and implementing the user interface, required communication modules, the data massage module, and database changes to accommodate the increased data requirements. Initially, a small group developed the system specifications and design. Subsequently the team was expanded to a peak of 14 people for coding, software installation and tuning, and system testing. This work was completed over a four month period. Finally, a smaller group made performance enhancements and supported the field roll-out. The total estimated effort of the Information Systems team was approximately 6 person-years.

The GE Corporate Research and Development (GECRD) team was led by a senior knowledge engineer (author Golibersuch) whose responsibilities included: knowledge engineering, development of machine learning tools, design and programming of the knowledge base, frequent communications with project leadership to assure that the evolving design was in accord with business objectives, and consultation on validation and system implementation. During prototype development he was assisted on a part time basis by several programmer trainees. During field testing of the pilot system he was joined by a full time programmer to re-engineer the knowledge base for production use, improve its performance, and organize the code for easy maintenance. The total effort of the GECRD team was approximately 1.5 person-years.

Knowledge Engineering

The most critical aspect of the entire project was to accurately capture the underwriting knowledge of experts and ensure consistency of that knowledge with risk management policy. Recognizing the critical nature of this process a full time expert was assigned and was supported by a committee of additional experts as described earlier. Most of the knowledge was captured and implemented using the Example Based Evidential Reasoning technique, which proved an effective tool for productive knowledge engineering.

Developing this knowledge was a multi-step process (see Figure 2):

1) The decision hierarchy was defined

2) Attributes for each node and possible values for each attribute were specified
3) For each node:
   a) The expert supplied examples and opinions usually accounting for 5-10% of the possible combinations for that node. These examples were entered into a commercial spreadsheet directly by the expert.
   b) The evidence values and interpretations of the combined evidence were derived. This derivation was aided by the learning algorithm implemented in a modified version of the spreadsheet program.
   c) The same spreadsheet program was used to "compute" an opinion for each of the expert supplied examples. The expert then reviewed any differences between these opinions and her own.
   d) The algorithm was then tested with "hold out" examples.
   e) The risk committee then reviewed the results and suggested changes and improvements.
4) Finally the complete knowledge base was tested with actual historical applications.

One of the many advantages of this process is that it identified potential problems well before proceeding to production. These include inadequacies such as: inappropriate hierarchy, missing or unnecessary attributes and values, inconsistent examples and opinions, conflicts with policy, and inappropriate recommendations. As such problems were identified, one or more steps in the above process were iterated until a consistent knowledge base was obtained that faithfully emulated the expert's behavior and met all business objectives and policies.

Implementation

System GENIUS was implemented within the existing GEMICO mainframe systems environment. This allowed for a seamless transition to the new approach without disruption to nationwide users. Also, much of the existing core software was reused.

As shown in Figure 4, the system was built using a modular architecture. Entirely new modules are the knowledge base "server" and the data massage module. Data massage converts the raw application data into a form suitable for use by the knowledge base. For example, this module computes total borrower income from various sources and ratios comparing this income to total debt and housing debt service. The user interface was reengineered to enable entry of the additional data required, for example, additional collateral information, and also to bring the various screens into alignment with the layout of the standard application forms used in the industry for improved data entry productivity.

When an underwriter calls up a particular application, the previously entered data are recovered and sent to the data massage module. The massaged data are then sent to the knowledge base to obtain the recommendation and supporting explanations. These are then returned to system control for display to the underwriter along with a brief application profile. This information is all contained on a single screen and it is the first screen seen by the underwriter. Subsequently the underwriter can proceed to disposition the application or obtain additional application details for further review. Depending on the situation, additional supporting data may be obtained or errors may be corrected. In such cases a new GENIUS opinion is obtained. The final disposition is recorded along with the most recent GENIUS opinion and explanation.

Commercial Tools Employed As a result of the previously discussed "Guidelines" project, GEMICO held a license for the KBMS shell product of Trinzic Corporation as well as experience installing and operating the shell within their mainframe environment. The constructs available within KBMS were judged adequate to program the knowledge base as designed. Another factor that favored this product was that GEMICO systems personnel, who primarily had backgrounds in traditional data processing, were comfortable with the KBMS syntax and Trinzic corporate "culture". These considerations led to selection of the KBMS shell. Most of the development and prototyping were done using the PC version and the code was subsequently moved to the mainframe version for production.

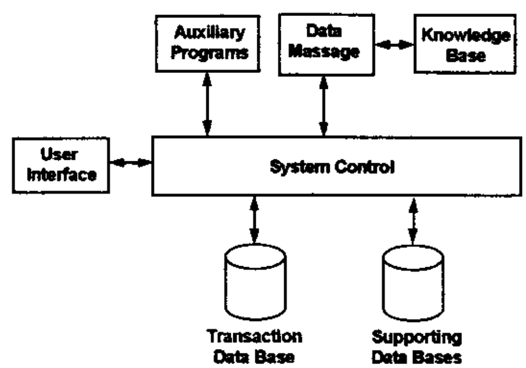

Figure 4. Mainframe system environment

The only other serious contender for use in the project was Inference Corporation's ART-IM. This shell was deemed to offer a much richer programming environment but suffered by way of compatibility with GEMICO's mainframe environment and culture. However, as will be further discussed under Future Plans, the ART-IM tool has been adopted to re-implement GENIUS as a PC based product for external customers as well as to support the business' move toward a distributed client-server environment.

The spreadsheet used was Microsoft Excel, which was both easy for the expert to use for direct entry and analysis of examples and sufficiently powerful to implement the learning algorithm. The only exception was a small 'C' program written to implement some of the required numerical computations. This was compiled as an external Excel function for superior performance.

Validation

Since the system would ultimately support underwriting of hundreds of thousands of applications, the potential consequences of inappropriate system operation or recommendations were quite serious. They included the possibility of substantial financial loss as well as damage to company reputation. As a result it was essential for a comprehensive validation of the completed system.

Validation of the knowledge base prototype as part of the knowledge engineering process described above was only the first step. Subsequent to integration of the knowledge base with the data massage, user interface, and other modules within the mainframe environment, the completed system was thoroughly tested. The underwriting expert, knowledge engineer, and systems personnel all participated in a thorough comparison of the recommendations of the system with prior field underwriter decisions with a large number of historical applications. Also, during this period a legal opinion of the system and basis for recommendations was obtained to ensure compliance with all laws and governmental regulations.

Subsequently the system was installed in a pilot branch. In a closely monitored situation, the system was tested by underwriters with actual current applications for a period of two months. From the time the system went into production all final GENIUS recommendations have been recorded and compared with the underwriter decision on an on-going basis. These data are periodically audited and reviewed to identify any problems or needs for corrections or enhancements of the knowledge base or other components of the system.

The importance of thorough and continual validation cannot be understated. As a result of the approach used in this project, problems with the GENIUS system in production have been few and minor. This has allowed all users—underwriters, risk managers, corporate managers, and ultimately customers—to rely on the system with confidence. Continuing acceptance is dependent on the most recent recommendation provided by the system. Therefore it is essential that the validation process be ongoing and especially important that it be repeated following any modifications to the knowledge base or any other part of the system.

Field Roll Out and Training

The system was introduced into the field offices on a gradual, regional basis. A plan was developed to ensure consistent use across geographical regions; provide appropriate, on-site training immediately prior to initial use; and minimize disruption of ongoing business. The plan took account of anticipated volume, business needs, and availability of personnel within the various regional offices. This approach ensured that each regional group of 1-5 branches, with the same regional management, would be converted to the new system within a short time period and thus encourage consistent use of the new system. This was particularly important because the old system continued to support areas of the country not yet converted. The entire roll out was completed in seven months. This includes the two month field validation in the pilot branch.

Minimal disruption of on-going business was a constant challenge as each branch's training and conversion were scheduled. Each branch was contacted well in advance so timing could be confirmed and preparatory work done. Training was done in half day sessions; half of the branch staff was trained in the morning and half in the afternoon. Training was supported by a manual, field by field explanations for data entry processors, and time for hands on practice. The business requirement to continue underwriting loans throughout the day of training was achieved without the need for sessions outside normal business hours. The following day (never a Monday) the branch was converted to the new system. The trainers remained on-site to support and assist the branch in adopting the new procedures. Due to the system's ease of use, the commitment to on-site training, and continued support during the first day of production use, the initial acceptance of the system was very good. Reliance on the system's recommendations was gradual over the following period of six to nine months.

Timeline

The project began in early 1992. Because of the prior work described earlier, the business requirements and technology options were well understood. Thus it was possible to assemble the team and proceed with development on a rapid pace. The knowledge base prototype was completed within nine months. The entire system was completed, tested, and installed in the pilot branch by early December 1992. The total elapsed time from project initiation to full scale production in June 1993 with approximately 4,000 transactions per day in 28 branches nationwide was 17 months.

Maintenance

The GENIUS system was designed to facilitate efficient maintenance. However, after two years of high volume production, the need for maintenance and thus our experience has been relatively limited. We believe the low maintenance to date confirms the careful knowledge engineering process and extensive validation of the original system.

All maintenance requires three steps:

a) Specification of the revised business knowledge or policies. The consequences of changes in business knowledge must be carefully considered before implementation in the system. Their impact will not be reflected in delinquency rate changes for many months and during that period many thousands of transactions will have been completed.

b) Implementation of the revised business knowledge via changes to one or more of knowledge base training, knowledge base and system code, and supporting databases. These changes have been made by the business' systems personnel after undergoing brief training on the design and operation of the knowledge base and EBER.

c) Revalidation. We have already stressed the importance of validation. The modular architecture of the knowledge base and system allows changes and thus most revalidation to be localized and focused. Prudence, however, dictates that for even minor changes the entire system should be revalidated using some form of regression testing to ensure that undesirable changes have not been introduced inadvertently.

Our limited experience indicates that steps 1 and 3 are the most expensive and time consuming.

The simplest and most frequent, indeed routine, form of maintenance is implemented by parameter changes in supporting databases. This is best illustrated by an example. The GENIUS system is in effect several systems whose analysis and recommendations are tailored for economic conditions. GEMICO continuously maintains a database so classifying all geographic areas of the country. Based on the location, the GENIUS system accesses this database to determine the economic condition appropriate for that application. Since the GENIUS system has been fully validated for each of these economic conditions from the outset there is no need for revalidation. Thus such changes can be implemented immediately upon determination that the local economy is changing.

So far there have been only two substantive changes to the knowledge base structure and training. The most extensive incorporated additional considerations to the assessment of revolving and installment credit. This required minor changes to the node structure, retraining of the credit nodes, and associated changes in supporting system modules such as data massage and entry screens. The second change was associated with refinement of the comparable analysis node that had been found to be excessively conservative. Such changes are straightforward and can be made in a time frame ranging from hours to a few weeks depending on complexity and any iterations required with business experts.

The generic explanation engine has worked very well and no significant changes have been made. Several additional exceptions to the generic engine were implemented to respond to underwriter requests that they be informed of certain attributes of a file even if the corresponding node has a favorable opinion. Such changes are easily made and require only limited regression testing to ensure that no inadvertent errors have been introduced.

There are several opportunities where the technical aspects of maintenance could be further automated. However, experience to date indicates that the overall structure and training of the knowledge base are quite stable and that frequently changing circumstances such as local market conditions are readily handled by the parameter based approach. Thus, motivation to further automate maintenance has been low.

Use and Payoff

Production Experience

In traditional practice an underwriter spends 15-45 minutes analyzing a loan and making a decision. The GENIUS system is able to evaluate the salient facts of an application and produce a recommendation and comprehensive explanation in a few seconds.

Use of GENIUS is fully integrated into the overall underwriting process. A recommendation is provided to the underwriter on every loan as the first step of the underwriting process. Since June 1993, all loans fully underwritten by GEMICO for insurance have been reviewed by the GENIUS system at least once. Whenever a modification is made to a loan file (to correct errors or add additional supporting information), a new recommendation is obtained and presented to the underwriter. This ensures that the GENIUS recommendation recorded when the file is dispositioned is based on the final set of data. On average, the GENIUS transaction is run about 2.5 times for each file underwritten. In the first two years roughly 2,000,000 transactions have been run to underwrite some 800,000 applications.

Agreement rates are monitored on a monthly basis by the audit staff with results distributed openly across all branches. All data concerns, agreement questions, or recommendations for knowledge changes are identified in these on-going reviews. The agreement rate is also used to identify the need for additional training and monitor the consistency of underwriting across the national markets and product lines.

When the GENIUS system was initially implemented, many underwriters were skeptical regarding its ability to understand the subtleties of underwriting. Over the past two years, confidence has increased to the point of heavy reliance on the system's recommendation and identification of concerns in the file. Improved productivity and customer service are clear indications of increased confidence.

Objectives Achieved

High Quality Decisions Underwriter agreement with the GENIUS system is monitored on a continuous basis and monthly and quarterly reports are prepared. This monitoring includes data entry integrity and completeness checks. An overall agreement rate of 90% was established in the first month of use and has been stable since that time (24 months). The agreement rate between GENIUS recommendations to approve and underwriter approval dispositions is higher at 98%. For high risk loans, the GENIUS system was engineered to be conservative, and agreement rates for those situations have been lower.

As the loans underwritten by the GENIUS system have begun to mature, it is becoming increasingly clear that the system does have the knowledge to correctly categorize loans by level of risk. This is supported by the fact that loans recommended by the GENIUS system for decline but which were accepted by the underwriter have a delinquency rate nearly four times that for loans recommended for approval. Inversely, loans with approve recommendations that have been insured by GEMICO are performing as expected for low risk loans.

Improved Customer Service During the period that the GENIUS system has been in production, turnaround time on insurance applications has improved. Needs for additional information on decisions and their rationale are communicated clearly and effectively to customers.

Underwriter Productivity It is estimated that a 20% improvement in productivity has been achieved so far. This has enabled better resource allocation to accommodate changes in business volume. Substantial further productivity improvements are anticipated (see Future Plans below).

Adaptability The GENIUS system has made it possible to quickly modify underwriting policy in response to changes in business and marketing strategy. Changes in the knowledge base and related system modules have been made at low resource cost and with no disruption of service.

Efficacious Development We believe the track record clearly supports the efficacy of the selected development approach. Development of a complex knowledge base that has achieved all objectives and required only minor modifications after two years of high volume production was completed in nine months.

Bonus Benefits and Future Plans

The outstanding success of the GENIUS system has led to additional strategic opportunities that were not anticipated at the time the project was initiated.

The reengineering efforts that started in 1992 are continuing and have become a way of life at GE Capital Mortgage Insurance. Now, changes are embraced and implemented in a more rapid manner. A major outcome of this ongoing effort has been the design and development of a National Processing Center to achieve additional improvements in productivity. The GENIUS system is a key building block of this center as it enables automated work flow assignments based on loan complexity rather than underwriter availability. This facilitates the optimal utilization of varying areas of underwriting experience and expertise. In conjunction with other technologies, the ability to manage work flow on a local, regional, and national basis is a reality.

Given the realized benefits of the GENIUS system, GE Capital Mortgage Insurance considered the possibility of making it available to its customers. Field tests with several customers demonstrated that with readily achieved modifications the GENIUS approach could be used to support other underwriting processes such as mortgage origination. As a result an external product version has been developed. The knowledge base was re implemented in a PC environment using ART-IM as the development shell. A module was added to evaluate loan compliance (Bynum et al. 1995) with secondary market requirements such as those of Fannie Mae and Freddie Mac. A companion statistical score is being developed to support the need for a quantitative estimate of default probability.

The new product was announced at the annual convention of the Mortgage Bankers Association in October 1994 and initial customer installations are anticipated for 1995. This new product offers additional income opportunity through license sales. Much more importantly, however, it offers GEMICO the opportunity to strengthen its business position through enhanced customer relations and facilitated product delivery. It also offers an exciting opportunity for AI technology to improve productivity throughout the mortgage industry.

Beyond mortgage services, the technology and knowledge engineering methodology developed during the

GENIUS project are being adopted for other financial services business. Applications to commercial credit analysis and underwriting of municipal bond credit enhancement are at an advanced stage of development. Other applications to product selection, deal structuring, and collateral evaluation are being explored.

Conclusions

Insights

We believe that the success of this project as well as many others provides overwhelming evidence of the practicality and major benefits to be accrued from application of AI technology within the financial services industry (IAAI 1989-1994). However, full realization of this potential requires a major investment of business resources. Indeed, the GENIUS experience was that business personnel had to invest more time than the AI technologists. Why is this so? First, the business is committing its "heart and soul" to a software implementation. Second, the knowledge engineering process inevitably leads to a thorough review and reconsideration of existing business practices. Third, extensive validation is essential. Fourth, effective implementation calls for "cultural" changes throughout the business. Fifth, during the course of development, new business opportunities are frequently identified.

The authors believe that this project is in many ways a model for the successful development and infusion of AI technology into a business. Several factors contributed to the successful evolution of the GENIUS project, which we believe are highly relevant to other similar projects. Thus we conclude this paper by citing these "lessons learned".

Lessons Learned

While many factors contribute to the success of a knowledge based application we believe the following five factors are absolutely crucial.

Top Management Support The GENIUS project benefited from the strong and unequivocal support of top management including the President of GEMICO and the CEO of GECMC. In particular, this support was reaffirmed at times of inevitable reversals in a project of this scope, complexity, and numerous technical and business challenges.

Committed Experts Major participation, preferably full time, by one of the business' best experts is essential. In addition, ready access to additional people with ancillary expertise is necessary. This requirement is probably the most "expensive" commitment the business must make upon undertaking such a venture.

User is King No system can be fully successful if it is not embraced by those who will be using it as part of their daily business lives. The needs of users should be solicited at the outset of the project, heard, and the system should faithfully respond to those needs. (As an interesting anecdote, the name "GENIUS" was the winning entry in a "name the expert" contest among users.) This process needs to continue throughout the production life of the system. Users must be supported by adequate training and reference material. In the most successful implementations the inevitable fears of displacement by new technology should be replaced by the perception of job enrichment and career enhancement.

Close Knit Team Developing and deploying a knowledge based system requires many talents and the best ideas of all. Managers, business experts, engineers and technologists, trainers, and even artists should all contribute in an environment that does not respect hierarchical or organizational boundaries. All disciplines must endeavor to clearly communicate their knowledge. At times this demands patience on the part of technologists to clearly explain to managers and business experts what the technology can *as well as* what it cannot do. Conversely, managers and business experts must take the time to clearly elucidate the business needs and processes to technologists. Disagreements and "family feuds" must be identified, aired, and resolved promptly without fear of retribution.

Careful Technology Choice It is necessary to carefully consider all potential technology options and make an objective selection of the most promising. As a corollary, if a choice is made that proves ineffective in meeting the business goals, it is imperative to set that selection aside and begin anew while preserving as much as possible the intellectual investment.

Acknowledgments

In addition to the authors, many individuals contributed to the successful development and deployment of the GENIUS system. The contributions of the following people are gratefully acknowledged. GEMICO Risk Management Team: Art Feingold, Debbie Johnson, Tom Hassinger, Joe McGuire, and Anne Hauser. User Task Force. Rosalie Peterson, Judy Johnston, Mike Mahoney, Brenda Richardson, Mary Walker, Monica Reid, and Mike Lepore. GEMICO Information Systems Team: Gail Morris, Anne Saleeby, Debra Williams, Chris Leu, Barbara Moresi, Debbie Mobley, Mac Davenport, Anand Komandur, Scott Short. Kay Conrad, Don Mills, Bob McKeown, Dale Smith, and Mary Phelps. GE Corporate Research and Development Team: Mary Phillips, Jayen Malde, Tim Smith, and Patrick Taylor.

References

Schneider, H. 1994. Great Strides in Industry Automation. *Mortgage Banking* 55/2:10-16.

Dennis, M. W. 1992. *Residential Mortgage Lending, Third Edition* (Chapter 8, Mortgage Insurance). Englewood Cliffs, NJ: Prentice-Hall, Inc.

Still, D. D.; Short, R. S.; Williams, D. B.; and Golibersuch, D. C. 1991. Guidelines: A Production Rule Based Implementation of Mortgage Insurance Underwriting Policy. In Proceedings of the First International Conference on Artificial Intelligence Applications on Wall Street, 174-177. Los Alamitos, CA: IEEE Computer Society.

Thomas, L. C.; Crook, J. N.; and Edelman, D. B. eds. 1992. *Credit Scoring and Credit Control*. Clarendon Press.

Talebzadeh, H.; Mandutianu, S.; and Winner, C. F. 1994. Countrywide Loan Underwriting Expert System. In Proceedings of the Sixth Innovative Applications of Artificial Intelligence Conference, 141-152. Menlo Park, CA: American Association for Artificial Intelligence.

Weiss, S. M. and Kulikowski, C. A. 1991. *Computer Systems That Learn*. San Mateo, CA: Morgan Kaufmann Publishers, Inc.

Golibersuch, D. C. 1995. Example Based Evidential Reasoning. Forthcoming.

Buchanan, B. G. and Shortliffe, E. H. 1985. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (Chapter 10, Uncertainty and Evidential Support). Reading, MA: Addison-Wesley Publishing Company.

IAAI 1989-1994. In Proceedings of the Innovative Applications of Artificial Intelligence Conference, Volumes 1-6. Menlo Park, CA: American Association for Artificial Intelligence.

Bynum, S; Noble, R.; Todd, C.; and Bloom, B. 1995. The GE Compliance Checker: A Generic Tool for Assessing Mortgage Loan Resale Requirements. In Proceedings of the Seventh Innovative Applications of Artificial Intelligence Conference. Forthcoming. Menlo Park, CA: American Association for Artificial Intelligence.