April 1990

ARTIFICIAL INTELLIGENCE APPLICATIONS IN LOGISTICS INFORMATION SYSTEMS

FINAL REPORT

Prepared for

Logistics Information Management Support System Program Office
Electronics System Division - AVSL
Air Force Systems Command
Hanscom AFB, MA 01731-5000

U.S. Department of Transportation
Research and Special Programs Administration

Prepared by

U.S. Department of Transportation
Transportation Systems Center
55 Broadway
Cambridge, MA 02142-1093

PREFACE

The Air Force Logistics Information Management Support System (LIMSS) program has prepared this report for the Air Force Logistics and Engineering (L&E) Information Systems Division (LE-I). This report is the principal deliverable from the LIMSS-AI project. It summarizes the results of a survey of existing applications and discusses the feasibility and benefits of specific candidate logistics applications.

The LIMSS Program Office staff at AFSC/ESD provided overall direction for this project. The work is being completed at the U.S. Department of Transportation, Transportation Systems Center (TSC), by Logica Technology Systems, Inc., with guidance and support from TSC LIMSS Program staff and from EG&G Dynatrend.

CONTENTS

EXECUTIVE SUMMARY

One of the responsibilities of the US Air Force Logistics and Engineering Information Systems Division (LE-I) is to identify ways in which emerging technologies can enhance the effectiveness of the logistics process. One technology which has in recent years provided significant benefits in a wide range of applications, including logistics, is artificial intelligence (AI).

In order to promote judicious use of AI within Air Force Logistics, LE-I has initiated an effort to identify AI applications with a significant payback. The task is being performed as part of the Logistics Information Management Support System (LIMSS) program sponsored by LE-I. The AI technologies considered are expert systems, natural language processing and advanced human computer interaction techniques such as intelligent front ends, adaptive user interfaces and intelligent training systems.

A survey of existing applications in logistics has highlighted a number of AI systems fielded, with more under development, in a range of logistics functions. Most activity to date has centered around expert systems, particularly stand-alone advisory systems and intelligent help for Air Force procedures and form-filling. Less activity has been identified in other strategically important technologies such as intelligent front ends to heterogeneous databases. The extent to which AI has been used also varies between logistics functional areas.

In identifying future candidates for AI systems, the aim has been to balance quantitative goals such as high payback applications with more qualitative ones such as the desire to encourage use of AI in functional areas where there has been little activity and the need to focus on strategically important but underused technologies such as intelligent front ends. A model for prioritizing projects has been developed that allows all these factors to be taken into consideration.

A number of candidate applications have been identified through interviews with potential users, government organizations responsible for developing applications and representatives of the LIMSS program. The prioritization model has been applied to the candidates and two of the highest ranking have been followed up in more detail. Others will be investigated in the near future.

As a result of the effort described in this report, four main activities are recommended as further work:

1) Establish functionality of a load planning expert system kernel. One means of improving return on investment is to develop a reuseable application kernel that can be applied to a range of similar applications. A load planning kernel would be of benefit in a range of transportation functions and a small study is required to specify the functionality, size and scope of such a system.

2) Investigate feasibility of an intelligent front end to LOGDIS. Intelligent front ends are a strategically important technology to the Air Force and other large organizations who wish to maximize the use that can be made of

existing investment in databases. An intelligent front end to the LOGDIS system would allow users straightforward access to a variety of databases and reduce the need for contract programmers. An assessment of possible architectures is required to establish the feasibility and likely cost of an intelligent front end.

3) Survey development methodologies. An important role for LE-I is to take a lead in issues such as methodologies for developing operational AI systems. A survey of existing work in AI system lifecycles and other areas of software engineering is a high priority in order to spread awareness within the Air Force of the applicability of available methods.

4) Continue application assessment. New candidate applications need to be prioritized and existing high priority applications followed up to a stage where the user organization can apply to existing funding sources. Critical issues in AI system deployment must continue to be highlighted, for example, the problems involved in providing accurate payback assessments.

Now that a network of AI contacts throughout the Air Force has been developed, LE-I is in an excellent position to act as a facilitator in the development of high priority applications. The activities that have been performed under this effort have filled an important need in encouraging useful and feasible applications. The four tasks recommended above will help take this promising first stage a step further towards the widespread use of AI to generate significant paybacks in Air Force logistics.

1 INTRODUCTION

One of the responsibilities of the US Air Force Logistics and Engineering Information Systems Division (LE-I) is to identify ways in which emerging technologies can enhance the effectiveness of the logistics process. One technology which has in recent years provided significant benefits to a wide range of applications, including logistics, is artificial intelligence (AI).

In order to promote judicious use of AI within Air Force Logistics, LE-I has initiated an effort to identify AI applications with a significant payback. The task is being performed as part of the Logistics Information Management Support System (LIMSS) program sponsored by LE-I.

In order to successfully identify high-payback logistics applications, two main tasks have been performed. The first surveyed existing applications of AI in logistics and other related fields to provide an overview of the kinds of application which can benefit most from the technology. The findings from this task were presented in the interim report "Artificial Intelligence Applications in Logistics Information Systems - Part 1: Survey". For completeness, these findings are summarized again in this report and updated as a result of ongoing activity.

The second task was to identify specific applications that fall within the framework of past successes and make an initial evaluation of the feasibility and potential paybacks of the most promising applications. The results of this task form the main focus of this report.

To ensure that a wide range of logistics functions are considered, the report has followed the practice of the LIMSS program and divided Air Force Logistics into a number of distinct functional areas:

- Supply
- Munitions
- Maintenance
- Transportation
- Engineering and Services
- Command and Control.

Section 2 of this report gives an outline of each of the AI technologies included in the analysis and Section 3 summarizes the findings from a survey of existing logistics applications. Section 4 discusses issues in assessing payback and presents the evaluation criteria that have been developed for this effort. Section 5 identifies candidate applications within each of the functional areas and presents the results of applying the evaluation criteria to each candidate. Section 6 continues by proposing follow-on activities to further develop the applications with the most potential and to continue tracking new candidates. Section 7 concludes with a summary of the main findings of the report.

2 ARTIFICIAL INTELLIGENCE TECHNOLOGIES

The AI technologies included in this task are expert systems, natural language processing and a class of applications arising from advanced human-computer interaction modelling techniques.

Other AI technologies such as speech understanding and neural networks have been considered for their potential in the longer term but have not formed the main focus of the task.

2.1 Expert Systems

Expert systems are a class of applications which model the way in which human experts solve a particular task. They are characterized by an explicit, declarative representation of the knowledge used in performing a task, rather than the implicit. procedural way in which knowledge is encoded using conventional programming languages. This means that the computer representation reflects more closely the concepts and approaches used by the human expert in solving a problem.

The human-like representation and reasoning used in expert systems lead to a number of functional advantages over conventional approaches, including user understanding of the system's problem solving processes, a high degree of user interaction at different stages of the task (for example, through 'what-if' analyses) and easier system maintenance. Organizational benefits can also result, such as improvement in the average quality of task execution, more consistent interpretation of policy throughout an organization and more time available for critical personnel to pursue other tasks.

In identifying potential expert system applications, it is important to consider the the various classes of problem to which an expert system could be applied as well as the full range of ways in which the system could support the user in performing a task.

Expert system applications generally fall into certain classes of problem to which the standard knowledge representation and inference techniques are well suited. Within the logistics arena, the relevant categories of expert systems are:

- diagnosis: for example, of equipment
- configuration: reducing the range of solutions to avoid a combinatorial explosion of possible configurations, for example, in planning loads for ships or planes
- planning and scheduling: broad range of applications that define an optimum way to organize future events to meet constraints, for example, in mission planning or policy formulation
- monitoring: constantly evaluating data and generating an alarm under specified conditions

- conventional applications: tasks which could be implemented by conventional means but which are developed using expert systems shells due to speed of development and ease of maintenance, such as advice on procedures and form-filling

- requirements specification: use of expert systems as a prototyping tool for defining organizational policy prior to implementation by conventional techniques.

The ways in which expert systems can assist a user in performing a task are varied. Early expert systems such as Mycin (Buchanan and Shortliffe, 1984) played the role of an independent expert that could be consulted much as a human expert could be. Issues of user acceptability and legal and moral responsibility soon led to a reevaluation of the role that an expert system should play within an organization. This, along with the natural creativity of people in exploiting new technology, has led to expert systems assuming a variety of roles with respect to their human users.

The independent model still exists, often as a fist line of consultation when a problem arises, For example, terminal or computer operators who have a problem with their equipment might turn to a diagnostic expert system as a first attempt to establish what the problem is. Recommendations from the expert system could be further diagnostic tests, corrective actions or instructions to call a human expert.

The most widely used role for an expert system is that of a decision support tool. A human expert engages the expert system in a dialogue to identify and evaluate alternative solutions before arriving at a final decision. For example, the Automated Air Load Planning System (AALPS) assists load planners in exploring alternative solutions to loading military equipment on board various types and numbers of military aircraft (Anderson and Ortiz, 1987).

As expert systems technology has matured, there have been moves to integrate expert systems with conventional technology. Embedded expert systems form part of a larger system in which different kinds of technology are used, each where most appropriate, in order to provide superior overall results. In the Pilot's Associate program, for example, a number of expert systems cooperate with each other and with other computing systems to provide efficient processing of cockpit data for interpretation by the pilot (Simpson, 1988).

In some cases it may be desirable to have the expert system play more of a background role. Hidden expert systems can lie behind applications and only alert the user to their presence when necessary, for example when an error or alarm condition is detected.

A more recent role for expert system is as a central repository of knowledge in a particular field. In developing an expert system, a large proportion of the time is spent collecting and modelling knowledge about the field. This may be drawn from existing information sources, such as regulations or textbooks, but more often involves encapsulating the previously undocumented experience of experts. By including knowledge browsing facilities such as entity-relationship diagrams and an information retrieval language, the knowledge within the expert system can be make available to users at any time, not just when they are performing a particular task (Lenat et al, 1983).

One of the current trends in expert system development is the use of application kernels to capture generalities within a class of problems. For example, a product formulation kernel has been developed (Skingle, 1987) which has been applied to a whole range of product formulation problems from pesticides to oil, and a computer-aided diagnostics system is on the market (Carnegie Group, 1989) that can be applied to a range of manufacturing problems.

The kernel has a core problem solving ability for the class of application which can be applied to each new specific application with a minimum of effort. This approach is becoming increasingly popular with AI system developers as a means of maximizing the reusability of and the investment in complex software. Where possible, Air Force AI projects should aim to adopt the kernel approach during architectural design.

2.2 Advanced Human Computer Interaction

Human-computer interaction is concerned with all aspects of a computer system and its users. In a move to improve the effectiveness and acceptability of computer systems, knowledge-based modelling techniques have been applied in a range of interactive applications. Without necessarily incorporating specific human task expertise, these applications use the explicit, declarative knowledge representation techniques used in AI systems to offer enhanced functionality to the user. Examples include intelligent front ends, adaptive user interfaces and intelligent computer-aided instruction.

Intelligent front ends are becoming increasingly important within the Air Force and other large organizations as the need to integrate heterogeneous systems grows. In order to make informed decisions, Air Force personnel often require access to information from a range of different databases and computer systems. Rather than rebuilding or duplicating information whenever a new requirement emerges, front ends are used to exploit existing investment in database development by enabling the same databases to be accessed by a range of users for a variety of purposes. Products that provide access to databases on different hardware platforms are becoming available, but even if the problems of physical communications are solved, an intelligent front end is required that can hide the details of the various database structures from the user and provide a coherent, integrated, task-oriented user interface.

As intelligent front ends are likely to be of strategic importance to the Air Force in the future, this effort has placed particular emphasis on identifying an application which will demonstrate their feasibility and potential.

Adaptive user interfaces operate by maintaining models of each user's capabilities. The models are used by the system to tailor information to users with different levels of expertise and also to each individual's changing requirements over a period of time. The models are updated by factors such as frequency of use and the amount of help requested. Adaptive user interfaces are particularly appropriate in systems which help a user to follow procedures; in these cases the level of detail of the guidance offered can be reduced as familiarity with the procedures increases. The benefits of an adaptive user interface include reducing the time taken to perform a task and enhancing user acceptability.

Intelligent training systems offer a mechanism to provide personnel new to a particular task with instruction tailored to their personal needs without the excessive costs of employing human training resources. Training can be offered in areas requiring either factual or procedural knowledge and can be supplemented by simulations where appropriate. Many intelligent training systems contain models of common errors and can therefore offer constructive advice in response to student errors. Some systems also contain a browsable model of the field under instruction, which encourages student exploration of the material in addition to the system-driven course.

2.3 Natural Language Processing

Natural language processing is a field concerned with computer analysis of natural human languages such as English, Spanish or Japanese. Originating with machine translation from one language into another, natural language processing is now becoming increasingly widespread for applications such as database queries and message understanding.

No natural language processing techniques available today can understand completely unconstrained natural language with all its ambiguities, specialist vocabularies, virtually limitless subject matter and ungrammatical yet understandable constructions. However, a number of systems have achieved a broad enough coverage to perform specific tasks in a useful manner. Some systems also have the ability to respond intelligently to non-understood sentences. These capabilities have led to the introduction of natural language processing into organizations such as the Coast Guard, the Navy and the Central Intelligence Agency.

The two main areas of natural language processing that are of most potential benefit within logistics are natural language interfaces to databases and message understanding.

Natural language interfaces to databases enable users to make database queries without the need to learn a formal query language. The natural language input may be a free-format request or it may involve building up a query through the selection of sentence fragments. Commercial products are currently available which use information about database structure to understand queries. Such systems require some tailoring to incorporate the terminology for each new database.

Message understanding systems can be used to summarize messages into a standard format or to route a message to an appropriate person or department. For example, a system developed for the Navy analyses ship-to-ship telex messages (Marsh et al., 1984), and a system built for a news agency routes new stories according to their topic (Hayes et al., 1988).

Message analysis systems have become increasingly widespread over the last few years and are generating commercial interest. Products are usually sold as application-specific kernels that need tailoring to a specific client's needs, such as ATRANS for banking telexes from Cognitive Systems. No commercially available message analysis systems have been identified for military use, although a number of prototype systems have been developed.

3 SUMMARY OF EXISTING APPLICATIONS

In order to decide on the potential for future applications, it is important to examine past success and failures in similar areas. Previous examples can be examined for lessons concerning technical feasibility, the role of a system within an organization and payback benefits. In addition, case histories provide inspiration to the user communities in seeing how the technologies could be of use in their own working environments.

In identifying existing systems within logistics, this analysis drew upon compiled summaries where possible. Two useful sources were the database of applications from the Air Force Logistics Command (AFLC) AI Support Center (AFLC, 1989), and the database of general DoD applications maintained by Simpson (1988). Specific system descriptions in journals and conference proceedings provided more detail, as did interviews with government agencies and contractors who are developing AI systems. The main emphasis was on Air Force logistics, although relevant systems developed by other agencies and organizations also provided useful insights.

Of the AI systems identified, the majority by far fall into the category of expert systems. This is in line with expectations, as expert systems are becoming an increasingly mature and widely accepted computing technique, while other AI fields such as natural language processing and intelligent front ends still have the potential for much broader use than at present.

Within expert systems, there are clear trends concerning the kinds of system that have been developed to date. In the AFLC database of applications, three types emerge as dominant in terms of numbers developed:

- stand-alone advisory expert systems using a question and answer session of the type provided by rule-based shells
- systems containing decision trees to provide guidance and training in following Air Force procedures
- error-checking expert systems which assist with filling in forms.

This emphasis has been due mainly to the policy of generating widespread awareness of AI through training in and use of rule-based shells on PCs. This has been highly successful in terms of education and propagation of small-scale systems, which collectively offer benefit to the Air Force. A similar strategy adopted by the company DuPont resulted in a return on investment on software and labor costs of 15:1 after two years of investment, even allowing for the initial steep learning curve (Feigenbaum et al, 1988).

Having established the utility of expert systems through small-scale developments, the AFLC AI Support Center now plans to focus on larger, integrated expert systems for maintenance and manufacturing.

To a large extent it is necessary to look outside the field of logistics in order to identify examples of successful natural language processing systems and advanced human computer interaction techniques. However, there are some notable exceptions, including an intelligent training system under development at the Logistics Management Center which is flexible enough to replace human trainers, and the ERIK natural language message analysis system that has been reading ship position reports for the Coast Guard since February 1985 (Cognitive Systems, 1989).

Within the six functional logistics areas, the amount of activity is varied. A number of small Supply expert systems for guidance with procedures and form filling are in place, while much less activity has been identified within Munitions. Maintenance expert systems are currently the subject of considerable interest and activity, although there are as yet few major systems fielded as they tend to be very large efforts involving integration with the whole equipment development lifecycle. Transportation has some diverse applications, including natural language message analysis and expert systems for load planning and transportation regulations. Engineering and Services is another area where there has been little activity, although a small number of systems have been identified that provide good examples of potential applications. Within Command and Control, the power of decision support expert systems is being applied effectively to planning and scheduling tasks.

Figures 3.1 through 3.6 summarize examples of existing systems in each functional area.

SUPPLY				
Name	**User**	**Status**	**Reference**	**Description**
Form filling	AFLC	Fielded	AFLC 1989	A number of applications which check that a form has been correctly filled out
IMA	AFLC/DLA	Fielded	Allen 1986	Assists an Inventory Management Specialist in validating data elements for use with D041
Warehouse training	AFLC	Evaluation	-	Computer based training system for warehousemen that has adaptive testing

Figure 3.1 Examples of AI Systems in Supply

MUNITIONS				
Name	**User**	**Status**	**Reference**	**Description**
KBLPS	Army	Prototype	HEL 1990	Tools for planning ammunition inventories and transportation
Reqs Forecasting	Army	Prototype	HEL 1990	Prototypes exploring learning algorithms for forecasting ammunition requirements

Figure 3.2 Examples of AI Systems in Munitions

MAINTENANCE				
Name	**User**	**Status**	**Reference**	**Description**
CEPS	SAC	Prototype	Papenhausen 1986	Helps B-1B technicians deal with the combinatorial explosion of fault diagnosis trees
ACES	Kelly AFB	Fielded	AFLC 1989	Assists help desk personnel diagnose data terminal and communications problems
IMIS	USAF	Prototype	Link et al. 1987	Interfaces with on-aircraft and ground-based systems to provide maintenance information
MITT	NASA	Evaluation	-	Trains flight coordinators and astronauts on the troubleshooting of NASA fuel cell

Figure 3.3 Examples of AI Systems in Maintenance

TRANSPORTATION				
Name	**User**	**Status**	**Reference**	**Description**
MACPlan	MAC	Fielded	Kissmeyer & Tallant 1989	Used by airlift planners to help in generating guidelines for aircraft and airfield use
AALPS	USAF	Evaluation	Anderson & Ortiz 1987	Assigns the loading of military equipment onto various types of aircraft
LOTS	Army	Prototype	Delfin Systems 1988	Helps an ocean terminal commander organize the off-loading and transfer of cargo
ERIK	Coast Guard	Fielded	Cognitive Systems 1989	Reads and interprets natural language ship position reports

Figure 3.4 Examples of AI Systems in Transportation

ENGINEERING AND SERVICES				
Name	**User**	**Status**	**Reference**	**Description**
Alarm Interpretation	Aeronautical Systems Div	Fielded	-	Interprets alarm conditions of a laboratory with environment control equipment
CCMAS	Engineering & Services	Prototype	Gregory 1990	Evaluates and manages construction costs for buildings

Figure 3.5 Examples of AI Systems in Engineering and Services

COMMAND AND CONTROL				
Name	**User**	**Status**	**Reference**	**Description**
FRESH	Navy	Fielded	Babin et al. 1987	Embedded expert system that evaluates alternatives in assigning resources to missions
WARPLAN	Personnel Readiness Ctr	Testing	Inference 1989	Assists in assigning qualified personnel resources for wartime/contingency operations

Figure 3.6 Examples of AI Systems in Command and Control

4 CRITERIA FOR EVALUATING PAYBACK

4.1 Issues in Payback Assessment

In the past, more attention has been focussed on establishing the technical feasibility of expert systems and other AI technologies than on the assessment of payback potential. For example, in a set of criteria for assessing expert system applications in Waterman (1986), the fact that "Task solution has a high payback" is just one criteria among many more technically-oriented ones such as "Experts agree on solutions" and "Task requires heuristic solutions".

As the number of successfully fielded AI applications increases, AI will begin to lose its special status as a technology warranting dedicated funding and will take its place as simply another tool for system developers to apply as appropriate. The result will be that more emphasis will be placed on justifying a potential application in terms of costs and benefits, as with any other planned system development. To a certain extent this is already taking place, as funding sources such as Productivity Investment Funding (PIF) evaluate proposals in terms of a minimum ratio of return on investment, regardless of the technology involved.

One problem in producing rigorous payback justifications is the range of ways in which AI systems can offer benefits. Reducing the time required to perform a task or preserving expertise that will be lost by retirement of key personnel are benefits that are often difficult to quantify. For example, the WARPLAN system developed for the Personnel Readiness Center at the Pentagon has reduced the planning time required from 2 weeks to 30 minutes. Payback calculations based on 2 weeks manpower savings do not accurately reflect the major operational impact that the system has during exercises or during a real crisis. Another example is in maintenance, where increased turnaround time for repairs can lead to an overall reduction in the number of units that need to be purchased.

Where savings can be quantified, they show that appropriate use of AI technology can give impressive results. The ERIK message analysis system has generated 90% cost savings for the Coast Guard, while the warehouse training system at the Logistics Management Center is expected to produce a return on investment of over 120:1. Findings from the commercial world are similar, with companies saving tens of millions of dollars per year and gaining returns on investment in the thousands of percent (Feigenbaum et al, 1988).

A closer analysis of existing systems that have generated a high payback show that they have a number of features in common that may be good indicators of the payback potential of planned systems. These features include:

- Stability of the task. Most of the tasks are fairly stable, building on well-established (although maybe unwritten) expertise and so avoiding high maintenance costs.
- Number of users. As with any system, a large number of users is a good indicator of a high return on investment.

- Cost of error. Even where smaller number of users are involved, a system that reduces the number of high-cost errors, for example in process control, can rapidly recoup development costs.

In order to test and refine these indicators of high payback, more information is needed regarding the actual return on investment figures that are being obtained, particularly in the Air Force environment. In many cases, the only return on investment data available are the initial estimates made in the application for funding; no analysis of actual achievement is being made once the project is completed.

In addition to immediate return on investment, it is also important to plan for the longer term. As discussed in Section 3, the spread of AI technology usage is not uniform among the logistics functional areas. In areas where there has been little activity, the initial emphasis should be to demonstrate low risk AI applications to generate user enthusiasm and understanding and management support of investment in the technology. Once this has been achieved, attention can turn to the major paybacks that are the longer term goal.

4.2 Evaluation Methodology

Candidate applications for AI projects are being collected through the following activities:

- interviews with representatives of the LIMSS program who are familiar with the functional areas
- telephone interviews with users and with organizations responsible for developing systems
- visits to user sites.

The size of the Air Force logistics community means that a small effort such as this cannot hope to obtain a comprehensive coverage of the users and so is necessarily opportunistic, but the involvement of the LIMSS representatives with their overview of user problems has helped to ensure that functional priorities are reflected in the applications considered.

Once a candidate has been identified, an initial screening is performed to ensure that AI is of relevance to the application. Sets of criteria for selecting expert system applications, such as the one in Waterman (1986), have been widely published and discussed and there is now general agreement about the types of application that are suitable for implementation as an expert system. For the other technologies, the principal selection criterion at this stage is that some kind of knowledge-based processing be required, such as models of expertise, tasks or users. For example, an intelligent front end will be included if it requires a model of where different kinds of data reside or knowledge about how to combine data for various types of task. The eventual implementation environment, for example database versus AI toolkit or C versus Lisp, is not viewed as critical factor.

Following this screening, each application is investigated in more detail to assess its potential benefit to the Air Force. According to a recent LE-I report (LE-I, 1989), there are two main aspects to prioritizing projects:

- the development of a set of criteria
- the establishment of a process for using these criteria to evaluate alternatives.

In developing the criteria for assessing the AI projects, a number of factors need to be balanced. LE-I, Air Force and DoD policy must be reflected in any decision, including the current emphasis on reductions in direct costs and manpower. Potential funding mechanisms for the proposed candidates are important and the majority of these demand a justification in terms of a high return on investment ratio. As a technology-oriented task, it is also a major goal to promote technologies that will be strategically important to the Air Force but that are currently under-used. Finally, project aims to encourage the use of AI in functional areas where there has been little activity in the past.

The following criteria have been selected to reflect the above project factors:

- The size of the system as an approximate indicator of development costs, as higher development costs lead to a reduced return on investment.
- Two of the indicators of high payback discussed in the last section: number of users and stability of task. The third indicator, reduced cost of error, is included as a type of benefit rather than as a separate criterion.
- An initial assessment of the types of benefit that the system would offer, categorized as:

  1) time saved through quicker decision (i.e. one person's ability to make more decisions or spend time on something else)

  2) manpower savings through less need for specialists

  3) reduced cost through better decision (e.g. higher availability of equipment)

  4) improved readiness through faster information or turnaround time (e.g. quicker plan development, faster diagnostics)

  5) possible improved readiness through better quality decision (i.e. identifying an opportunity that might otherwise have been missed).
- The type of technology, with more weight placed upon strategically important but under-used technology.
- The functional area, with more importance being given to areas of low activity.

Criteria	Values	Numeric Rating
Size	Large (>$1 million)	1
	Medium ($150K - $1 million)	2
	Small (<$150K)	3
Number of Users	1-99	1
	100-999	2
	1000-9999	3
Stability	Low	1
	Medium	2
	High	3
Benefits	Quicker Decision	1
	Improved Readiness (Quality)	1
	Manpower Savings	2
	Reduced Cost	3
	Improved Readiness (Speed)	3
Functional Area	Supply	1
	Maintenance	1
	Transportation	2
	Command and Control	2
	Munitions	2
	Engineering and Services	3
Technology	Expert System	1
	Adaptive User Interface	1
	Natural Language	2
	Intelligent Training	3
	Intelligent Front End	3

Figure 4.1 Criteria Values with Numeric Rating

A scoring model is used to apply these criteria in evaluating the candidate applications. First of all a numerical rating is assigned to each possible value within a criterion, with a higher rating indicating a higher priority application. Following this, the criteria themselves are weighted for their relative importance to the goals of the effort. Figure 4.1 shows the values and corresponding numerical rating for the criteria and Figure 4.2 gives the weighting of each criterion. Finally, each candidate is assessed against the criteria and an overall numerical indicator of priority calculated by:

- assigning the candidate a value for each criteria

Criteria	Weighting
Size	1
Number of Users	5
Stability	3
Benefits	5
Functional Area	5
Technology	3

Figure 4.2 Relative Weighting of Criteria

- substituting the numeric rating equivalents for criteria values
- multiplying the numeric rating by the weight assigned to each criterion
- summing the weighted scores for all criteria.

The results of applying the scoring model to specific candidates are discussed in Section 5.7.

The benefits of this scoring approach to project prioritization include the use of both qualitative and quantitative criteria and the ability to provide an audit trail of decisions (LE-I, 1989). The main drawback is the subjectivity involved in assigning numeric rating equivalents for the qualitative criteria values. In this project, it is important to include the more qualitative factors such as importance of functional area and strategic implications of the technology, so the benefit of including these outweighs the drawback of the subjective numeric assignments. Similarly, the relative weighting of the criteria is necessary to ensure that the project priorities are properly incorporated into the final score for a candidate. Even if other interested parties do not agree with the precise numeric ratings for values or with the weighting of the criteria, the fact that all of these figures are explicitly stated makes it a straightforward matter to change the priorities and recalculate the results.

One of the most controversial sets of numeric assignments is for the expected benefits of the system. The numeric ratings for the 'Benefits' criterion shown in Figure 4.1 have been designed to incorporate Air Force policy, tempered by what can realistically be expected from AI systems. So, although a quicker response on the part of the decision-maker can be expected to generate some cost savings, such benefits are generally less significant that the effect that improved turnaround time can have on readiness. Similarly, AI can conceivably lead to manpower reductions, but these savings are likely to be smaller than those generated from reduced costs through better decisions, for example regarding inventory reductions. Improved readiness can be achieved in two ways: through faster decisions or through identifying an opportunity that would otherwise have been missed. Faster decisions are ranked higher than 'better quality' decisions, as in the logistics context the system will only be able improve upon human decision making on occasion, whereas faster decisions can consistently improve the turnaround time for action.

As the project moves forward, it is expected that the numeric ratings for criteria values, the criteria weighting and maybe even the criteria themselves will be modified to reflect changes in emphasis in LE-I policy or to compensate for the fact that more activity is taking place within a functional area or a technology.

Following the selection and prioritization of candidate applications, those that emerge as the highest priority are pursued to the third stage of analysis. The aim of this stage is to look at the candidate in enough detail to identify the actions and timescales necessary to develop the application further. For smaller applications, the recommendations may be in the form of a development plan and return on investment analysis for submission to a funding source such as PIF. For larger applications, the proposal is more likely to be for a feasibility study and prototype that will lead to a full development plan and costing.

Whether performed as part of this effort or as a separate feasibility study, the development of a plan involves a preliminary assessment of the targeted users, the tasks they perform and the functionality of a system that could assist them. In estimating the development costs, a consistent model such as the one described in Melaragno and Allen (1989) should be applied. Payback assessment is more of a problem, as models which quantify all relevant benefits are not readily available. More guidance is needed on how to perform this part of the return on investment calculation; it would be most beneficial if a standard method were developed and distributed widely throughout the Air Force.

The deliverables from this stage should be an outline description of functionality, an initial technical solution, a development plan including effort and timescales and a return on investment analysis.

5 PROPOSED APPLICATIONS

An initial screening of candidates led to ten applications being selected for further investigation. These are described in the following Sections 5.1 to 5.6, grouped according to functional area. Section 5.7 gives a comparison of the candidates according to the evaluation methods presented in the last section.

It should be emphasized that the applications presented in this section in no way represent a thorough analysis of Air Force functional needs and priorities. To obtain results in a short amount of time, the effort needed to draw upon the experience of those who have already assessed potential improvements within a functional area. In accepting the following applications for consideration, the assumption has been made that these candidates have been put forward because they offer more significant benefits than other candidates within the same part of the Air Force. These benefits will of course be validated before substantial effort is invested in an application. This assumption does not preclude the existence of equally strong candidates in other parts of the Air Force. Section 6.4 recommends an ongoing task to identify these other applications.

5.1 Supply

Supply is an information- and regulation-intensive activity that gives rise to many suggestions for expert support in collating large amounts of data, following appropriate regulations and completing forms in an accurate manner. Four of the applications that have been suggested to date are as follows:

1) The Cataloging and Standardization Center (CASC) at Battle Creek, MI, may soon be merging its Catalogers and Standardization Technicians into a single Logistics Data Manager designation. Traditionally, cataloging has been more clerical in nature, while standardization requires a technical specialist for each type of equipment. The organizational change will lead to changes in the skills profile of the staff performing cataloging and standardization, with a need for technical support in some of the equipment types. Expert systems which provide causal models of equipment operations or which help make key decisions about interchangeability and substitutability would ease the problems caused by the changes.

   Each expert system would be used by between 1 and 10 personnel. The systems would range in size between small and large, depending on the complexity of the equipment being modelled, but the most useful would tend to be on the large side. Once developed for a piece of equipment, the system would be fairly stable and should not require major modifications. The benefits of developing the system would be a more accurate assessment of interchangeability and substitutability possibilities through an improved knowledge of the workings of the equipment, which would result in improved readiness. The system may also alleviate the need to recruit staff with very specialized technical knowledge.

Expert systems which model the operation of complex equipment are also in demand in the maintenance area. In the longer term, the development of diagnostic expert systems as part of the main equipment development lifecycle may lead to the availability of expert models which could also be used at CASC. However, in the short term it is worth investigating how specific equipment models might help maintain a high quality of decision making at CASC in the face of the organizational changes.

Contact: Sandy Schauder, CASC/SCX

2) Through the Standard Base Supply System (SBSS), information on supply support levels is available to base-level personnel. The interpretation of this information is not always straightforward, particularly identifying the relationship of supply support levels to weapon system support. An expert system approach to trend analysis would complement some of the statistical approaches already available and planned for the future, and would allow users who are not familiar with the workings of the statistical approaches to perform some of the required analysis.

   The system would be used by approximately 800 base-level staff and would be a medium-sized system. This would include the effort required for the development of the statistical analysis and graphical displays as well as for the expert system. There may be changes required to the system over time as new performance measures are developed and integrated. The benefits associated with use of the system would be in identifying opportunities for inventory (and hence cost) reduction that might otherwise have been missed and in improving readiness by maintaining more appropriate stockage levels.

   Contact: Col William Shaw/Major Kathy Johnson, HQ USAF/LEYS

3) The Logistics Management Center has identified a need for a system to assist in ordering items of equipment. There are a large number of regulations governing how to place an order, and mistakes are often made in preparing the required forms. This results in delays as incorrect order forms are returned for resubmission. The situation is further affected by the high turnover of staff.

   An expert system containing Tables of Allowances and other regulations and providing context-sensitive advice based on the information entered so far would avoid a significant amount of wasted time and effort. The system would be moderately large, due to the hundreds of tables involved, but once developed it would be used by over 3000 base-level personnel. The system would not need to be updated very often as the regulations only change every few years as new technology evolves. The main benefit of the system would be in improving the turnaround time of orders by avoiding multiple resubmissions, which could have an overall impact on readiness.

   Contact: Lt Col Peterson, AFLMC/LGK

4) The information needed by item managers may be located in a number of different databases. Although the item managers have access to all the necessary databases, they are not able to see more than one database on the screen at one time. Front end software that displayed the databases in multiple windows, or a more intelligent system that collected the information and integrated it, would allow far more rapid comparisons between data.

A simple front end which used multiple windows would be a fairly small system, but would not really involve artificial intelligence technology. An intelligent front end which performed specific comparisons would involve a medium-sized system. The task is relatively stable and the expected number of item managers who would use the system is 5000. Savings would result from improving the speed of decisions by an item manager.

Contact: Mike Riley, LMSC/SYU

5.2 Munitions

Problems for which AI could provide a solution in the field of munitions include storage optimization and asset realignment. One candidate that has been proposed is as follows:

5) A system for munition storage optimization, including repositioning when a new item arrives, could be used by base level personnel. If a generic load planning kernel were developed (see Section 6.1), then munitions storage could be a specific application produced using the kernel.

The system would be medium sized and offer a service to 1800 users. The techniques used for optimization are relatively stable; a system for performing repositioning is currently under development and could form the basis of the algorithms. An expert system could help the planner explore alternative scenarios more easily than the system currently under development and would lead to time savings in the planning cycle.

Contact: Wayne Benson, HQ USAF/LEYWN

5.3 Maintenance

As discussed in Section 3, expert systems applications in Maintenance are currently receiving considerable attention from other AI programs in the Air Force and for this reason have not been the main topic of this effort. However, maintenance troubleshooting is an extremely important application for expert systems and the following is typical of the many candidates within Maintenance:

6) In maintenance troubleshooting, technicians are confronted with large volumes of fault diagnosis trees, many of which are found to be inaccurate on close examination. A computerized system would provide a more manageable medium for following the relevant trees and could help uncover inaccuracies in the existing trees. The Air Force Logistics Management

Center sponsored the development of a system for fault reporting, but the problem of fault isolation is much larger. A rule-based approach is not likely to be appropriate, so the LMC is interested in pursuing model-based approaches.

The system would be specific to a particular weapon system, for example the F15, and would be used by over 10,000 maintenance technicians and pilots. The development of a full model of the weapon system would be a large undertaking, although once developed the knowledge would not change very frequently. This is an example of an application with a high direct cost of error; it is vital that faulty components are identified rapidly and accurately. Benefits of the system would be a quicker decision and improved readiness due to faster and better identification of faults. In the longer term, a faster turnaround time for repairs could lead to a reduced inventory and lower costs.

Contact: Lt Col Duvall, AFLMC/LGM

5.4 Transportation

Problems in Transportation center around the planning of optimal routing and use of resources. One application that occurs in a number of similar areas is that of load planning, where an embarkation officer is responsible for producing an optimal load plan for a given vehicle with a given set of cargo. Similar issues arise whether the vehicle being loaded is a truck, aircraft, ship, warehouse or other container, and the following proposal seeks to capture this commonality.

7) Determining an optimal load plan for a vehicle is a time consuming task. The planner needs to consider constraints such as the dimensions of the loading space, the weight and size of the cargo, regulations concerning the stowage of hazardous cargo, interactions between types of cargo and speed of loading and unloading. The process often requires a number of iterations before a satisfactory plan is generated.

An expert system kernel which captured the common elements of load planning and allowed rapid tailoring to a specific application could improve the speed of decision making for a wide range of potential users. Assuming that the kernel could be adapted to loading planes, trucks, trains and ships, the user base would be in the region of 300. This could be increased if the kernel could also be applied to warehouse storage of supplies and munitions. There is also the potential to apply the kernel to load planning in the other services, which would increase the number of users still further.

The techniques used in load planning have been developed over many years of experience are not usually subject to much change. The initial development of the kernel would be a medium-sized effort, and tailoring it to each specific application would require a small to medium-sized project. The major benefit lies in the time saved in producing a plan, not only in terms of manpower savings but more importantly in terms of the turnaround time for reacting to circumstances during a crisis. The system may also be able provide a more thorough check of all the regulations and other constraints than a human could.

This application is discussed further in Section 6.1.

5.5 Engineering and Services

Engineering and Services covers a very wide range of applications. Because of this diversity the potential for AI applications is large, but the same diversity has made it difficult to gain an overview of where the highest paybacks are likely to be. The following proposals represent two candidates that are currently being pursued.

8) As with weapons systems maintenance, engineering maintenance is an area where AI can be of significant benefit. There are a large number of different services that need maintaining, such as heating, ventilating and air conditioning (HVAC) equipment, fire alarm systems and generators. An expert system for troubleshooting would provide support during training and actual operations. Given the large number of similar potential applications, it is important to select one for an early demonstration of feasibility in this area. HVAC has been identified as one of the highest priorities for improved troubleshooting.

   The system would be used by about 6 personnel at each of the 100 bases, a total of 600. The system would be of medium size and high stability, although there may be some tailoring required at each base to reflect local conditions. As in most maintenance applications, the cost of error can be high and the benefits arise from faster and more accurate diagnosis with the corresponding gains in cost reduction and readiness.

   Contact: Larry Strother, AFESC/DEM

9) A topical issue that is causing problems within Engineering and Services is that of privatization. Staff turnover makes it difficult to maintain the knowledge about legislative authorities, legal and public issues and economic analysis that is required for informed decisions. An expert system which gave support in some or all of these areas would both save time and also improve the final quality of the decisions, which can involve up to $150-250 million.

   The system would be used extensively by 20-30 Engineering and Services staff, with another 150 to 200 making occasional use. The system would need to be large to cover all the requisite background for decision making, although a number of smaller systems that dealt with specific difficult areas would also be beneficial. The knowledge in the system would be subject to change as the regulations were updated.

   Contact: Al Nixon, AFESC/DEQ

5.6 Command and Control

Command and Control is an area that can benefit significantly from the powerful decision-making environments that expert systems can offer. One of the key aspects of Command and Control is the integration and evaluation of data from different sources. The following candidate demonstrates how AI can form part of an intelligent front end for decision-makers to access multiple databases.

10) The LOGDIS system offers Air Logistics Centers (ALCs) and AFLC HQ access to a number of databases to support a range of decisions. However, in order to convert the data into a form that is useful for decision-making, the users must either manually integrate the data or wait up to six months for a purpose-built script to be built. An intelligent front end is needed which would support a wide range of requests and integrate the data into a useable format, but which would not require much training to use. Natural language would be a highly desirable interface medium.

    The front end would be used by approximately 5000 ALC and AFLC HQ personnel. The system would be large, but once developed would be moderately stable, subject only to modifications in the structure of the databases it accesses. It would lead to quicker decision making and, for some types of decision, improved readiness as a result of the reduced turnaround time.

    Contact: Jim Henley/Mike Riley, LMSC/SYU

5.7 Evaluation of Candidates

The evaluation methodology presented in Section 4.2 was applied to the candidates described above. Figure 5.1 shows the values for the criteria that were assigned to each candidate and Figure 5.2 show the total score calculated by substituting the numerical equivalents and multiplying by the criteria weighting.

Both of the maintenance diagnostics expert systems appear in the top ranked applications due to the significant benefits that this kind of system can bring, while the LOGDIS intelligent front end scores highly because of the strategically important technologies involved. The privatization and the load plan kernel expert systems represent useful applications with reasonable benefits and good scores in most of the criteria.

Those high scoring applications which were identified early on in the project have been the subject of more detailed investigation, the results of which appear in the next section. Others were not proposed until more recently and will form the main focus of future analysis. In parallel with the more detailed analysis, the prioritized list of applications will be maintained through the addition of new candidates.

Application Name	Size	Users	Stability	Technology	Area	Benefits
1) CASC Standardization	Large	1-10	High	ES	Supply	Manpower Savings Improved Readiness/Quality
2) SBSS	Med	800	Med	ES	Supply	Reduced Cost Improved Readiness/Quality
3) Supply Ordering	Med	3000	High	ES	Supply	Improved Readiness/Speed
4) Item Manager IFE	Med	5000	High	IFE	Supply	Quicker Decision
5) Munitions Storage	Med	1800	High	ES	Munitions	Quicker Decision
6) Fault Diagnosis	Large	10000	High	ES	Maint.	Quicker Decision Reduced Cost Improved Readiness/Speed
7) Load Plan Kernel	Med	300	High	ES	Transpt.	Quicker Decision Improved Readiness/Speed
8) E&S Maintenance	Med	600	High	ES	E&S	Quicker Decision Reduced Cost Improved Readiness/Speed
9) Privatization	Large	150	Med	ES	E&S	Quicker Decision Reduced Cost
10) LOGDIS IFE	Large	5000	Med	IFE NL	C&C	Quicker Decision Improved Readiness/Speed

Figure 5.1 Criteria Values Assigned to Candidate Applications

Application Candidate	Total Score
8) E&S Maintenance	74
6) Fault Diagnosis	68
10) LOGDIS IFE	67
9) Privatization	55
7) Load Plan Kernel	54
3) Supply Ordering	49
2) SBSS	46
4) Item Manager IFE	45
5) Munitions Storage	44
1) CASC Standardization	38

Figure 5.2 Candidate Applications Ranked by Total Score

6 SELECTED APPLICATIONS

One problem in obtaining funding for new system developments is that a significant amount of effort may be required to scope out the problem before any return on investment calculations can be performed, but funding will only be granted once these figures are available. One of the aims of the current LE-I initiative is to alleviate situations like this by working with users to scope applications with good potential prior to submitting a request for funding.

Section 5 reported on some of the candidates that have been proposed to date. Two of the highest priority candidates have been selected for further investigation and are described in more detail in this section. For each candidate, further discussion of the application and outline plans for the next stage of development are given.

In addition to these specific applications, recommendations for work in two other areas are also presented. The first involves methodological issues in fielding AI systems and the second is a continuation of the current effort to survey and identify AI applications and issues.

6.1 Load Planning Expert System Kernel

The load planning expert system kernel has been selected to demonstrate how expert system kernels can help increase the payback potential of software through reuse of the kernel in different parts of the Air Force.

An expert system kernel is a means of capitalizing on the commonality running through a class of applications. Load planning is performed in a number of functional areas in the Air Force (and other DoD organizations), for example in Transportation for aircraft loading, in Supply and Munitions for safe storage and in Command and Control for personnel and equipment loading. Capturing the common elements of load planning in an application kernel would increase the number of users of the system and so improve the potential payback.

Load planning applications have a number of common features, such as:

- regulations governing distances between objects
- the need to separate certain kinds of materials from each other
- weight and dimension considerations
- unloading priorities
- unloading routes, including vehicle turning circles etc.

A common approach to load planning could be captured in an application kernel expert system that offers the appropriate knowledge representation and reasoning techniques for these applications. The kernel could then be tailored with minimum effort to incorporate specific regulations, requirements and priorities. Figure 6.1 shows the kind of architecture that would be appropriate for a load plan kernel.

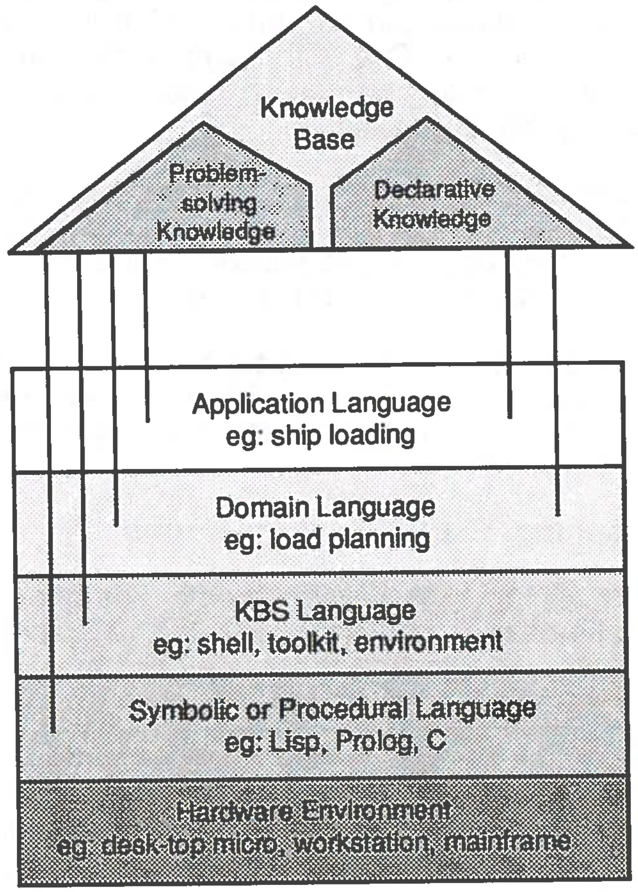

Figure 6.1 Architecture for a Load Plan Expert System Kernel

In producing an expert system kernel, it is advisable to begin by developing two specific applications. These applications should not be too similar too each other; they should contain features covering the range of anticipated applications for the kernel. As the knowledge acquisition, functional specification and initial design proceed, those elements which are common across the two applications are factored out into a separate component of the system, which becomes the kernel. Once the two applications and the kernel have been fully developed, the kernel needs to be tested on a third application. The development of the third and subsequent applications using the kernel should be significantly faster than the initial development. Insights gained from each new application may lead to minor modifications to the kernel, but if the first two applications have been carefully selected, then the majority of requirements will already have been covered.

To progress the idea of an expert system for load planning further, the first stage is to assess more accurately the size and scope of the system and its potential paybacks. This information can then be used to apply for funds for the development activity. The activities required in this first stage include:

1) identifying two or more specific applications that contain a range of the features required in the kernel and gaining strong support and commitment from the user, expert and management representatives

2) broad survey of all potential target applications which could use the kernel to identify total number of users, projected paybacks and target hardware and software environments

3) initial knowledge acquisition from expert load planners in the two application areas to identify the types of reasoning and constraints used

4) initial attempts at factoring out common reasoning and other knowledge to demonstrate feasibility of the kernel approach

5) initial outline of functionality for the applications and for the kernel

6) development of a small prototype which performs one aspect of load planning using kernel knowledge tailored to each of the applications

7) production of a plan for full development, including a cost justification for obtaining funding and recommendations for the hardware and software environment, the development lifecycle, the task breakdown and the effort required.

The development of a prototype is not essential at this stage, but it does help considerably in obtaining support for further development. If it does not form part of this first stage, then a prototype should be developed in the early phases of the development activities as a mechanism for early feedback on functionality, user interface and design issues.

This first stage would take approximately 6 man-months of effort, including protoyping. Gaining the support of appropriate user representatives is a prerequisite for any further progress and so should be completed before commitment is made to the rest of the effort.

The deliverables from this stage would be a report describing the results of Tasks 1 to 5, prototype software from Task 6 (if included) and a separate plan for a full development project as described in Task 7.

6.2 LOGDIS Intelligent Front End

The LOGDIS intelligent front end is an excellent example of how AI techniques can be applied to enhance the usability of a computer system other than through an expert system. As well as providing significant benefits in its own right, this application would demonstrate the feasibility of developing an intelligent front end to heterogeneous databases, a technology which is of strategic importance throughout the Air Force and DoD.

The Intelligent Gateway Processor (IGP) subsystem of the Logistics Data Integration System (LOGDIS) offers users at Air Logistics Centers and at AFLC HQ access to multiple databases on a variety of machines. However, in order to draw on the available information, users either have to perform multiple system interrogations and then integrate the retrieved data manually, or they have to wait for purpose-built Network Access Machine (NAM) scripts to be developed by contract programmers, which can take up to six months.

An intelligent front end to the IGP could significantly enhance the utility of LOGDIS. A possible architecture could involve the following components:

- **Database Index**: A model of the location and format of data available through LOGDIS. This would act as an index to the databases for use by the other modules. The Command Dictionary/Directory (CDD) currently being developed by AFLC may provide the basis for this module. Information about the number of records in the databases would also be useful in alerting the user to an expensive search before it is actually entered.

- **User Interface**: A user interface which accepts English sentences and uses a natural language processing system and the Database Index to formulate a series of queries in SQL and other appropriate query languages. The English sentences could be typed in directly or could be created through the selection of query fragments from menus. The benefits of the latter approach are that the user can see the kind of data that is available in the system and can only generate queries that the system is able to understand.

- **NAM Script Generator**: Writes or selects the NAM scripts necessary to make the queries created by the User Interface. This would involve selecting from a library of existing scripts to log on to the various remote hosts and selecting or generating scripts to make the queries once connected.

- **Results Integrator**: Once the information has been retrieved from the remote databases, the results need to be integrated into the form requested by the user and displayed in an appropriate manner.

This is just one possible approach to the design of a front end. Before arriving at a final design, other approaches need to be investigated. For example, a tool which helped the user build their own NAM scripts and queries through menu selection and on-line help might be a simple and effective option. A more detailed assessment also needs to be made of the kinds of queries that users generate most frequently. For example, if a large proportion of the queries involve retrieving data from just one database, then the system could link the user directly to that database and carry out queries through a commercial natural language interface such as NLQ from Battelle.

The following tasks are recommended to further investigate the feasibility of the LOGDIS intelligent front end:

1) analysis of query types made by a representative sample of intended users

2) investigation of alternative architectural designs and their impact on end user functionality

3) discussion of functionality options with LOGDIS team and with representative users, resulting in decision on best option

4) development of a small prototype to demonstrate user functionality and feasibility of architecture

5) production of final report describing recommended functionality and architecture, proposed development plan for full system, estimated costs and cost justification.

As with the load planning expert system kernel, development of a prototype should ideally be included in this first stage to assess feasibility prior to committing funds to full scale development. Useful results could again be achieved with approximately 6 man-months of effort.

6.3 Development Methodologies

The third area of proposed future activity focuses on methodological issues of developing AI systems for operational use. In recent years there has been an increased awareness of the need for methodologies reflecting the specific characteristics of AI systems. AI developers and researchers are investigating issues such as system lifecycles, verification and validation, highly-interactive user interfaces, real-time performance, delivery vehicles and maintenance. Achieving results in these areas will affect the long term viability of fielding AI technology. The Air Force must maintain an awareness of the guidelines and standards that are emerging and ensure that they are applied in an appropriate manner within logistics AI projects.

A number of organizations are developing methodologies that include issues in AI system development. For example:

- Boehm (1988) has developed a spiral model for software development which incorporates the prototyping activity often associated with AI projects
- the IEEE Computer Society has recently established a series of committees as part of the Expert Systems Subgroup
- the Rome Air Development Center (RADC) has a major project to provide software engineering support for the system development lifecycle
- the Nuclear Regulatory Committee is investigating issues of verification and validation of safety-critical expert systems

- in Europe, collaborative efforts such as the GEMINI project sponsored by the UK Government's Central Computer and Telecommunications Agency encourage companies involved in developing AI systems to work together to define a commonly acceptable methodology.

Given the amount of existing activity, there is little point in LE-I sponsoring the development of yet more guidelines and methodologies. However, it is important that developers of AI systems within the Air Force are aware of what is available and are encouraged to apply proven techniques rather than creating a new set for each new project. LE-I has a role to play in reviewing and evaluating methodologies that have been developed or are under development and in recommending when the methodologies should be applied.

The main purpose of an LE-I effort in the area of methodologies would be to carry out a survey of existing methodologies and guidelines and widely circulate recommendations as to their appropriate usage. This would involve the following activities:

1) literature and telephone survey of organizations which are active in the development of methodologies and guidelines

2) telephone survey of AI system developers within and outside DoD to identify current practices in the development and delivery of operational systems

3) in-depth analysis of 3-4 case studies of the methodology applied by organizations in developing and fielding a successful system

4) preparation of a report documenting the findings of Tasks 1-3

5) preparation of a concise report summarizing the methodologies that have proven most effective and recommending when and how these should be applied by AI system developers in Air Force logistics.

A thorough and useful analysis of methodologies and the creation of the resulting reports could be performed in 6 to 8 man-months. Once the main effort had been completed, the information could be reviewed periodically to incorporate new developments in the field as part of the ongoing task described in Section 6.4.

6.4 Continued Application Assessment

In a project to identify potential applications, at any particular point in time some candidates will have been developed to a greater level of detail than others. The candidate applications and proposals for further work described in this document represent the status of the project as of April 1990.

In parallel with the preparation of this report, the project is continuing to identify new candidates and track and develop existing ones through discussions with user representatives. The aim of these activities is to make an informal assessment of the technical feasibility and an initial evaluation of the potential value to the Air Force. For those candidates which look the most promising, some effort will be invested in discussions with the users in order to mature the proposal to a stage where the user organization is in a good position to apply for development funding. The project has no influence over the actual selection of projects for funding.

It is important for a policy-level organization such as LE-I to encourage the development of strategically important projects within the Air Force. Selected candidates should reflect current priorities and policies, for example by:

- incorporating under-used technology that will be of future importance, such as intelligent front ends to heterogeneous databases
- demonstrating the feasibility of applying AI in functional areas where there has not been much activity to date, such as Engineering and Services and Munitions
- developing expert system application kernels that can be applied across Air Force functions and even across services.

Maintaining a central view of activities throughout the Air Force is also critical in order to avoid duplication of effort. The survey of existing applications needs to continue by maintaining contact with developers of AI systems for logistics and updating the database of applications in each functional area as new projects are identified.

In addition to identifying new candidates and surveying existing applications, a third role for LE-I lies in ensuring that the AI systems being fielded are effective and reliable. Issues such as methodologies for AI system development need to be addressed centrally and recommendations circulated widely amongst those who are responsible for development and fielding.

In order to avoid conflicts with other Air Force departments, the role that LE-I plays in supporting the use of AI should not duplicate the activities of other AI groups. The main central organization for AI in Air Force logistics is the AFLC AI Support Center at Wright Patterson Air Force Base in Dayton, Ohio. This Center was established to promote the use of AI within AFLC and has been very successful in generating widespread awareness of the technology. Their activities center around the development of AI applications, tool and technology assessment and continuing education, all of which are extremely complementary to the role proposed for LE-I. Continuing interaction between the groups is important to coordinate activities; for example the AI Support Center is currently focussing on systems to support the maintenance of weapons systems, so the LE-I efforts can place a lower priority on this functional area and concentrate on others where there is less activity.

A number of other groups, such as the Logistics Management Center and the Aeronautical Systems Division are also involved with AI, but their main emphasis is on developing specific AI applications, which is not not included in the LE-I brief.

To support the required LE-I activities in AI, a small ongoing effort is required. The current model, with one person working approximately half time and input as needed from other LIMSS representatives, is working well and could be continued. The following activities should be included in any further effort:

- assistance in progressing the projects described in Sections 6.1 through 6.3
- following up in more detail the most promising application candidates, currently the Engineering and Services maintenance and privatization proposals
- providing input to the LIMSS documents describing long term requirements
- continuing to identify and pursue applications, survey existing developments and investigate policy issues such as AI development methodologies and payback assessment.

Interim reports describing particular applications should be generated as appropriate, with progress reviews after 3 and 6 months. The 6 month review should involve preparation of a report summarizing progress and making recommendations for any further effort.

7 CONCLUSIONS AND RECOMMENDATIONS

The survey of existing logistics systems has shown that AI systems are in use in a range of functional areas within the Air Force and are beginning to offer significant benefits in cost reduction and improved operations. In particular, the better understanding in recent years of the power and constraints of expert systems has led to their wider acceptance as a computing tool for solving certain classes of problem. The main emphasis in the Air Force to date has been on stand-alone advisory expert systems and intelligent help for procedures and form-filling, while future plans involve increased emphasis on Maintenance expert systems integrated with other technology such as Automatic Test Equipment.

However, the AI successes identified so far are isolated examples and there is still outstanding potential for AI systems to offer substantial benefits in many areas of Air Force logistics activity. The applications recommended in this report are a first step in helping to realize this potential, but there are still more ways in which logistics operations can benefit from AI technology. For example, the power of expert systems to act as highly interactive decision support tools is not being fully exploited; more widespread use of the kinds of decision support applications being developed in Command and Control could produce significant operational improvements. Expanding the scope of the AI technologies considered to include natural language processing and the kind of advanced human-computer interaction modelling used in intelligent front ends will open up the range of applications even further.

One of the problems in selecting candidates for future activity lies in prioritizing the applications that are proposed by the functional representatives. Automation in general and AI in particular are likely to lead to improvements in most cases, but funds need to be distributed in an equitable manner to obtain maximum return on investment for the Air Force. The model that has been described in this report offers a mechanism for including both quantitative criteria that indicate a high payback application as well as qualitative ones that reflect other goals of the project. The scoring method itself has worked well, particularly in terms of the ability to incorporate priorities by weighting the criteria and to maintain visibility of the factors that lead to the conclusions.

The area of the model that needs more attention is in ensuring that the criteria used are really the best indicators of a high payback application. Case studies describing the substantial commercial savings resulting from AI have been publicized over the last two years (largely through the work of Feigenbaum et al., 1988), but within Air Force logistics there is little information available concerning paybacks from past applications. Based on the information that is available, a number of features of successful applications are emerging, but far more quantitative data is needed to fully characterize the types of application that are producing the highest paybacks.

Applying the prioritization model to a number of applications has produced a ranked list of candidates. Proposals for further development of two candidates - a load planning expert system kernel and an intelligent front end to the LOGDIS system - have already been outlined and others will be generated in the near future for the remaining high ranking candidates.

In addition to pursuing specific applications, an important role for LE-I is in building awareness within the Air force of issues which affect all AI system developments. For example, one recommendation of this report is to perform an analysis of system development methodologies and develop guidelines regarding the applicability of different approaches to the design, implementation and fielding of operational systems. Future efforts in this area could include recommendations on the use of models for return on investment and cost/benefit analysis.

The activities that have been performed under this effort have filled an important need in encouraging useful and feasible applications. LE-I should continue the activity of identifying and evaluating candidates and helping to develop high priority applications. The four tasks recommended in Section 6 will help take this promising first stage a step further towards the widespread use of AI to generate significant paybacks in Air Force logistics.

APPENDIX A

INDEX OF ACRONYMS

AALPS	Automated Air Load Planning System
AFLC	Air Force Logistics Command
AI	Artificial Intelligence
ALC	Air Logistics Center
CASC	Cataloging and Standardization Center
CCMAS	Construction Cost Management Analysis System
CDD	Command Dictionary/Directory
CEPS	Central Integrated Test System (CITS) Expert Parameter System
DLA	Defense Logistics Agency
FRESH	Force Requirements Expert System
HEL	Army Human Engineering Laboratory
HVAC	Heating, Ventilating and Air Conditioning
IGP	Intelligent Gateway Processor
IMIS	Integrated Maintenance Information System
IMA	Inventory Manager's Assistant
KBLPS	Knowledge-based Logistics Planning Shell
L&E	Logistics and Engineering
LE-I	US Air Force Logistics and Engineering Information Systems Division
LIMSS	Logistics Information Management Support System
LOGDIS	Logistics Data Integration System
LOTS	Logistics Over-the-Shore Planner
MAC	Military Airlift command
MITT	Microcomputer Intelligence for Technical Training
NAM	Network Access Machine
PIF	Productivity Investment Funding
RADC	Rome Air Development Center
SAC	Strategic Air Command
SBSS	Standard Base Supply System
TSC	Transportation Systems Center

APPENDIX B

ORGANIZATIONS CONTACTED

The following organizations have been contacted in connection with this effort:

7th Communications Group, Washington DC
AFLC Aeronautical Systems Division, Wright Patterson AFB OH
AFLC AI Support Center, Wright Patterson AFB OH
AFLC Cataloging and Standardization Center, Battle Creek, MI
AFLC Human Resources Laboratory, Wright Patterson AFB OH
AFLC Logistics Management Center, Gunter AFB, AL
AFLC LMSC/SYU, Wright Patterson AFB OH
AICorp, Waltham MA
Air Training Command, Randolph TX
Army Human Engineering Laboratory, Aberdeen MD
Army Log Center, Fort Lee VA
Carnegie Group Inc, Pittsburgh PA
Cognitive Systems Inc, New Haven CT
DARPA, Arlington VA
Delfin Systems, Sunnyvale CA
Engineering and Services HQ, Tyndall AFB FL
Gold Hill Computers, Cambridge MA
HQ USAF/LEYWN, Washington DC
HQ USAF/LEYS, Washington DC
Inference Corporation, Los Angeles CA
Intellogistics Inc, Columbus OH
Logica Cambridge Ltd, Cambridge UK
Neuron Data, Palo Alto CA
Quintus, Mountain View CA
Standard Systems Center Future Systems Office, Gunter AFB AL
Strategic Air Command, Offutt AFB NE
Systems Command Human Resources Laboratory, Brooks AFB TX
Systems Command HQ, Andrews AFB MD
SRI International, Menlo Park, CA
Tactical Air Command, Langley AFB VA

APPENDIX C

REFERENCES

AFLC, Database of AFLC AI Applications maintained by the AI Support Center of the Air Logistics Division (ALD/LTI), 1989.

Allen, M., The Development of an Artificial Intelligence System for Inventory Management, Council of Logistics Management, Oak Brook, IL, 1986.

Anderson, D. and Ortiz, C.,"AALPS: A Knowledge-based System for Aircraft Loading", IEEE Expert, Winter 1987.

Babin, M., Gately, M. and Sullivan, M., "FRESH: A Naval Scheduling System", Proceedings of the IEEE Conference on Artificial Intelligence Applications, 1987.

Boehm, B.W., "A Spiral Model of Software Development and Enhancement", Computer, May 1988.

Buchanan, B. G, and Shortliffe, E.H., Rule Based Expert Systems, Addison Wesley, 1984.

Carnegie Group, Publicity Brochure, 1989.

Cognitive Systems, Publicity Brochure, 1989.

Delfin Systems, "Delfin Systems", Defense Electronics, June 1988.

Feigenbaum, E., McCorduck, P. and Nii, H.P., The Rise of the Expert Company, Times Books, 1988.

Gregory, R.A., "Construction Cost Management Analysis System", AF Engineering and Services (AFESC/DEC), 1990.

Hayes, P.J., Knecht, L.E. and Cellio, M.J., "A News Story Categorization System", 2nd Conference on Applied Natural Language Processing, 1988.

Inference Corporation, "Expert Systems at the Pentagon", 1989.

HEL, Project Descriptions from US Army Human Engineering Laboratory, 1990.

Kissmeyer, K.Y. and Tallant, A.M., "MACPlan: A Mixed-initiative Approach to Airlift Planning", in Schorr and Rappaport, 1989.

LE-I, Development of a Prioritization Process for Information System Projects, LE-I Report, August 1989.

Lenat, D.B., Borning, A., McDonald, D., Taylor, C. and Weyer, S., "Building Expert Systems with Encyclopedic Knowledge", International Joint Conference on Artificial Intelligence, 1983.

Link, W.R., Von Holle, J.C. and Mason, D., Integrated Maintenance Information System (IMIS): A Maintenance Information Delivery Concept, Technical Paper, Air Force Human Resources Laboratory, 1987.

Marsh, E., Hamburger, H. and Grishman, R., "A Production Rule System for Message Summarization", National Conference on Artificial Intelligence, AAAI, 1984.

Melaragno, J. and Allen, M.K., "A Plan for the Application of Artificial Intelligence to D0D Logistics, Logistics Management Institute Report PL816R1,1989.

Papenhausen, D., "B-1B Maintenance Diagnostics Gain a Long Memory", Defense Electronics, September 1986.

Schorr, H. and Rappaport, A. (eds), Innovative Applications of Artificial Intelligence, MIT Press, 1989.

Simpson, R.L., Jr., "DoD Applications of Artificial Intelligence: Successes and Prospects", Applications of Artificial Intelligence IV, International Society for Optical Engineering (SPIE), 1988.

Skingle, B., Product Formulation Expert System, Report to Alvey Directorate IKBS/052, 1987.

Waterman, D.A., A Guide to Expert Systems, Chapter 11, Addison Wesley, 1986.