Great Lakes Fishery Commission
Project Completion Report *

DEVELOPMENT OF AN EXPERT SYSTEM FOR INTEGRATED MANAGEMENT OF SEA LAMPREY: A FEASIBILITY STUDY

Draft Report

by

M.L. Jones
J.F. Koonce
C.H.R. Wedeles

Submitted to
Great Lakes Fishery Commission
October 2, 1987

ACKNOWLEDGEMENTS

We would like to thank the staff of the Sault Ste. Marie Sea Lamprey Control Centre, particularly Mssrs. Stan Dustin, Jerry Weise, and Rod McDonald for kindly providing the data necessary to initialize the stream database. As well, data for the IMSL simulation model were provided by Bill Dentry (OMNR) and Cliff Schneider (NYDEC).

Special thanks also go to Linda Rattie for once again doing an excellent job of drafting the figures, editing, and producing this report.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

1.1 Background

In recent years a tremendous surge of interest in the potential use of a variety of technologies collectively known as artificial intelligence has emerged. One of the strongest areas of development in this large field is in expert or consultative systems, which seek to capture the expertise that exists within a particular discipline in a manner that allows individuals to directly access that expertise in support of their own decision-making. In recognition of this important new research and development area, the Great Lakes Fisheries Commission decided to fund a feasibility study of the potential use of expert systems for Great Lakes fisheries management. They chose as their subject the development of a prototype expert system for integrated management of sea lamprey, based on a proposal submitted by ESSA Environmental and Social Systems Analysts Ltd. and Dr. Joe Koonce of Case Western Reserve University. This report presents the results of the feasibility study.

Before describing the objectives of the study, it is worth briefly explaining the purpose and general structure of expert systems. The essential purpose of an expert system is to translate the knowledge of one or a number of experts into a symbolic form that includes the logic by which such experts would arrive at conclusions about questions lying within their area of expertise. If the expertise is properly represented, then others not having the same degree of expertise can use the expert system, rather than the expert himself, to assist their decision-making. Expert systems are especially useful, therefore, in situations where expertise is not widely available, or is distributed amongst individuals who are not easily consulted all at the same time.

A good example of an expert system is a tool for medical diagnosis. The body of knowledge available for diagnosing ailments based on their symptoms is now sufficiently large that it is very difficult for individual GP's to have all the necessary facts at their disposal. Expert systems have been developed that can be used by doctors, or even patients, to diagnose their symptoms. By querying the user for information on symptoms, the system follows a logical sequence to arrive at one or a number of suggestions as to what the problem might be.

Expert systems work by connecting a set of facts, some supplied by the user, others potentially within a database, to a set of rules by which information represented in the facts can be used to reach conclusions. Generally speaking, expert systems are used to ask two types of questions. First, one can ask, "What will happen if I do X?". Alternatively, the question could be, "How is it possible for Y to be true?". Using the medical diagnosis example, the "what if" question might be used to define the symptoms of a particular disease, whereas the "how possible" question would be used to identify which diseases might give rise to the symptoms described by the doctor or patient. It is clear from this example that in the former case the user would be working forward through the logical structure of the expert system while in the latter, one is working backwards.

In resource management situations, simulation models are often used to ask "what if" questions. The AEAM models developed in the early 1980's by the Great Lakes Fishery Commission are examples. Similarly, optimization techniques are used to ask "how possible" questions. Simulation models and optimization techniques are actually specific types of expert systems which almost always rely exclusively on quantitative information. The important difference between these two tools and expert systems in general is that the latter can and usually do make use of non-quantitative information, typically represented in the form of expert judgement.

Finally, expert systems offer one other important advantage as a decision support tool: documentation. Because of the explicit logical structure of an expert system, it is relatively easy to document the logic used by the system to reach a particular conclusion. An often cited problem in resource management (and other disciplines) is the lack of accountability for decisions that are made. Although expert systems should (and will) never replace decision making itself, at least they can be used to supply an explicit, logical rationale for decisions that are made.

1.2 Objectives

The specific objectives of the study funded by GLFC and conducted by ESSA Ltd. and Dr. Koonce were to:

1. select a case study for the development of a prototype expert system and identify the information needs (and sources) to provide the knowledge base for the expert system;
2. develop the prototype expert system using whatever information and knowledge is readily available;
3. present the prototype expert system to a small number of representatives of the Great Lakes fisheries management community; and
4. prepare and submit a brief report that documents the development of the expert system and identifies future information and development needs.

At a project planning meeting held in Ann Arbor on June 10, 1987, the project steering committee agreed that the topic for the case study would be Integrated Management of Sea Lamprey (IMSL). This choice was considered desirable for four reasons:

1. It is obviously a subject area of considerable interest to the Commission;
2. it is typical of most resource management situations in that the decision maker is faced with a moderately large number of management choices;
3. it is an area that lends itself to the use of both non-quantitative (expert judgements) and quantitative (simulation models) tools; and
4. it lends itself to the development of an explanation facility that documents the logic used to reach a given conclusion.

The remainder of this report is divided into three sections. In Section 2 we describe the prototype expert system. In Section 3 we illustrate its use and present a preliminary evaluation. Section 4 is devoted to the presentation of a series of recommendations for further development of the expert system based on discussions with the project steering committee at the expert system demonstration meeting (Objective 3). We have also included, as an Appendix, a complete listing of the expert system software developed during this study.

2.0 DESCRIPTION OF THE PROTOTYPE EXPERT SYSTEM

Having agreed that IMSL was an appropriate overall topic, it was then necessary to decide on the specific focus for the prototype development work. All members of the steering committee agreed that the expert system should not be restricted to the issue of maximizing the effectiveness of sea lamprey control per se, but rather should include consideration of the implications of control strategies in the context of broader fisheries management goals, particularly lake trout rehabilitation. Specifically, we decided that the expert system should be designed to assist the user in examining the consequences of various combinations of lamprey control and fisheries management strategies in terms of their effects on various indicators of system response. Thus, our emphasis was on allowing the user of the expert system to ask "What if" questions, as defined in Section 1 of this report.

2.1 Scope

The definition of the scope, or bounds, of the expert system determines which elements need to be included and where there are opportunities for simplification. At the June 10 planning meeting, the steering committee discussed the desired bounds of the prototype expert system. Specifically, we considered four major bounding issues: spatial context, temporal context, management goal measures, and management actions.

2.1.1 Spatial Context

For any resource management situation, one needs to define the area affected by management decisions. In the case of sea lamprey management on the Great Lakes, this area is clearly the Great Lakes basin. However, to simplify the problem for the purposes of developing a prototype, we chose to restrict the spatial scale to that of Lake Ontario and its tributary streams. In addition to having relatively few lamprey producing streams, the existence of an IMSL simulation model for Lake Ontario strongly influenced our choice.

Within this overall area there are distinct units of management that may or may not require explicit recognition in the expert system. In the lake itself, we decided that no spatial resolution was required and the parasite and host populations in the lake therefore could be treated as a single unit. In contrast, it was felt that the explicit treatment of individual lamprey producing streams would be a highly desirable element of an IMSL expert system.

2.1.2 Temporal Context

To be able to explore the consequences of a particular strategy in terms of its effect on indicators such as lake trout populations, it is necessary to have a reasonably long time horizon. We agreed that a minimum time scale would be approximately 15 years, a scale that corresponds to current fisheries management planning horizons. This means that projections generated by the expert system should extend at least 15 years into the future.

Within this 15 year time horizon, decisions are made on an annual basis. For sea lamprey control, there is a three year planning cycle that affects decisions taken in the subsequent three years. We decided that the expert system should be designed to function on an annual, or at least a triennial decision cycle.

2.1.3 Management Goal Measures

To judge the consequences of the management strategies chosen and implemented by the expert system, it is necessary to to identify a number of measures of success or failure. We identified ten quantitative indicators of the success of a lamprey control/fisheries management strategy:

- # parasitic phase lamprey
- # spawning phase lamprey
- # producing streams

- carcass density
- wounding rates
- lake trout harvests
- # spawning lake trout
- # wild yearling lake trout
- mean age of lake trout
- lake trout growth rates

Each of these indicators is an output of the existing IMSL model and thus should be easy to produce if the model and the expert system can be linked.

Two other management goal measures that were recognized as being important are the actual treatment schedule for the lamprey producing streams and the costs of lamprey control.

2.1.4 Management Actions

Finally, to develop and implement a management strategy the expert system must include those management actions that comprise the strategy. For the purposes of developing the prototype, we identified three major management actions to be included: chemical treatment of ammocetes, stocking of salmonids, and fish harvest controls.

Overall, the IMSL expert system is designed to examine the consequences of different combinations of stream treatment, stocking, and harvest control in Lake Ontario over a 15 year time horizon. The consequences are described in terms of changes over the 15 year period in ten lamprey and lake trout population parameters. The costs of sea lamprey control provide an additional measure of performance.

2.2 Expert System Structure

The prototype expert system comprises six major components (Figure 1). These six components sequentially allow the user to review and modify the initial knowledge base if new facts are available, develop a stream treatment strategy, implement that strategy using the IMSL simulation model, document the decisions taken by the user, and examine the results of implementing the strategy by looking at

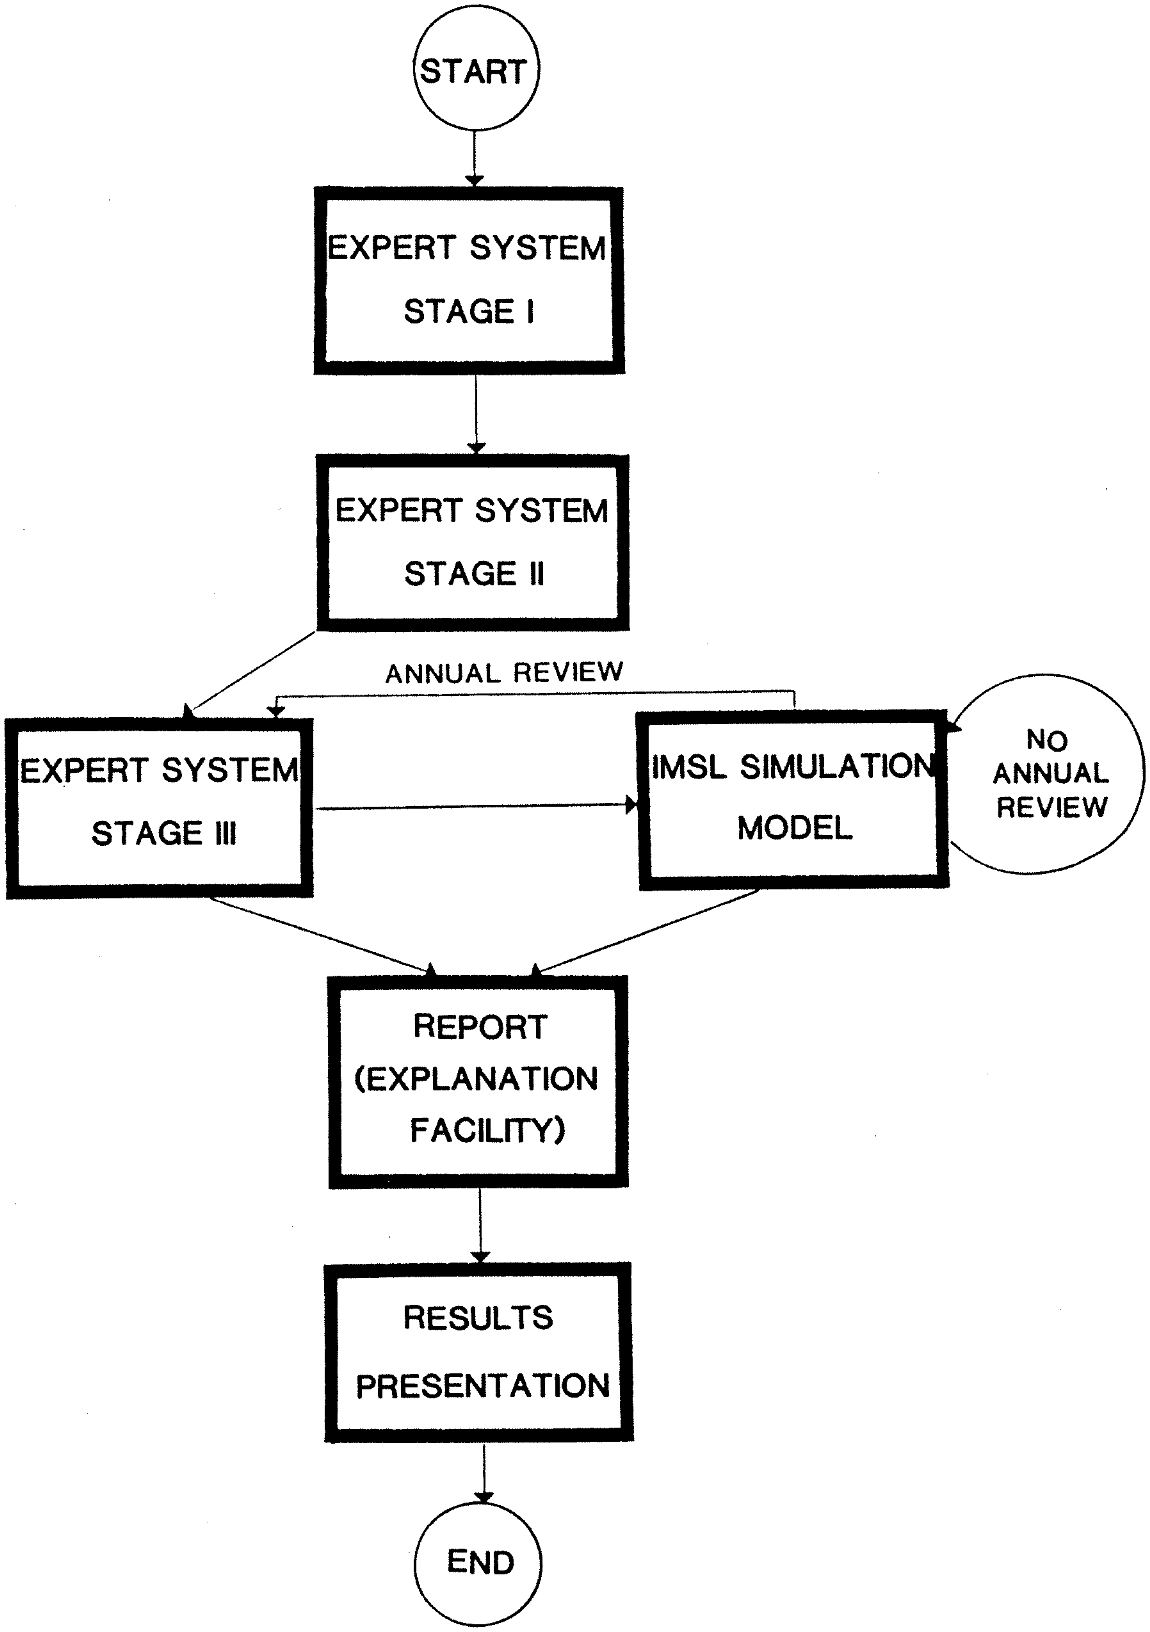

Figure 1: Overall structure of the prototype expert system.

changes in the management goal measures listed above. In the sections to follow we describe each of these components.

2.2.1 Expert System Stage I: Knowledge Base Update

The first component of the expert system is a simple database management facility. The knowledge base for the system essentially consists of a database containing a set of current (1986) conditions for each of the 49 potential candidate streams for treatment on Lake Ontario. The user is asked whether he would like to modify any of the database elements for any of the 49 streams, presumably because more recent information has become available. If the user chooses to update the database, then the expert system queries the user for the changes he would like to make.

The database contains information on eight attributes for each stream (Table 1). The first five of these are self-explanatory. The "effectiveness of stream survey" attribute was intended to reflect any independent information that might suggest the survey data itself underestimates ammocete densities, relative to other streams. An "historical production index" was coined to represent the possibility that some streams might be treated regardless of surveyed ammocete densities, simply because they have historically always produced large numbers of transformers. Finally, the "external risk rating" attribute was intended to account for streams where treatment would not be possible due to external factors, such as the presence of highly valued trout spawning runs in the stream at the appropriate time for treatment.

The attributes listed in Table 1 probably do not adequately represent the information used by control agents to judge which streams to treat. To develop a more comprehensive and representative database for an operational expert system would require a great deal more interaction between the experts (i.e., the control unit specialists) and the knowledge engineers (i.e., the programmers that develop the

Table 1: Stream database attributes.

Attribute	
• date of last treatment	for each stream
• chemical requirements	
• treatment area	
• density of ammocetes > 125 mm	
• ammocete age structure	
• effectiveness of stream survey	
• historical production index	
• external risk rating	

expert system) than was possible in this study. Nevertheless we feel the prototype database is adequate for the purposes of this feasibility study.

2.2.2 Expert System Stage II: Stream Treatment Strategy

The expert system next asks the user a series of questions that lead to a choice of stream treatment strategy. This component actually has two parts: choosing the criteria for stream selection (Table 2) and selecting an algorithm for stream ranking (Table 3).

Of the eight strea- screening criteria listed in Table 2, only four are included in the prototype. Two of these four criteria, historical production indices and external risks of treatment are, however, effectively ignored. In the absence of information required to initialize the stream database for these two criteria, all streams were assigned the same value. Finally, the size of the stream is implicitly included as a criterion in the stream ranking procedure, at least when costs are used to select streams (see below).

The user is asked which of these criteria he would like to include to select streams as candidates for treatment. Depending on which criteria are chosen, the user is then prompted for further information necessary to apply each criterion. For example, if "years since last treatment" is chosen, then the user needs to supply a value for the length of the minimum time period between treatments.

The application of these criteria allows one to obtain a list of streams that are legitimate candidates for treatment, according to the strategy chosen. It may not be possible (or necessary) to treat all of the streams on the candidate list in a given year. Rules are therefore required to determine which streams from the candidate list will be treated. These rules have two elements: a method of ranking all candidate streams and a criterion for deciding when enough streams have been treated (a stopping rule).

Table 2: Stream screening criteria.

• years since last treatment	**
• ammocete density > 125 mm	**
• ammocete size structure	
• historical production	*
• survey effectiveness	
• external risk of treatment	*
• size of stream	***
• constraints due to timing	

** included

* included, but not properly initialized

*** included implicitly as cost factor

Table 3: Stream ranking algorithms.

1. Cost/benefit with fixed budget.

2. Maximum benefit with fixed budget.

3. Cost/benefit with target ammocete reduction.

4. Maximum benefit with target ammocete reduction.

Table 3 lists the four combinations of stream ranking and stopping rules that are currently implemented in the prototype expert system. The user is asked to select one of these algorithms in the final step of Stage II. The cost/benefit algorithms rank streams according to the ratio of benefits to costs of treatment; the maximum benefit method ignores costs in ranking streams. The fixed budget stopping rule simply means that streams are treated in order of rank until the budget (user supplied) is exhausted. The target ammocete reduction stopping rule operates by treating streams in order of rank until a lake-wide target level of ammocete control (actually transformer production control) is achieved. Note that the cost/benefit with fixed budget algorithm, when used with years since last treatment and ammocete densities greater than 125mm, produces a stream treatment strategy identical to the backpack method of Heimbuch and Youngs.

2.2.3 Expert System Stage III: Strategy Implementation

Having obtained (from the user) the necessary information to determine a stream treatment strategy, the expert system then reads the initial stream database and applies the strategy to select a set of streams in the current year (presently set at 1986). Note that the database contains all the necessary information to implement the strategy.

Now the system is ready to begin to look at the consequences of implementing a particular strategy using the IMSL simulation model. The user is asked to supply a time horizon for the model (i.e., the number of years over which the model will be run). Then he is given the option of allowing the simulation to run for the specified period of time with no intervention, or being advised of the stream treatment schedule for each year of the simulation and being given the opportunity to add or delete streams from the schedule. After making these choices, the system then executes the simulation model over the chosen time horizon.

2.2.4 IMSL Simulation Model

The IMSL Simulation Model is essentially the same as that described in Spangler and Jacobson (1985) and adapted for Lake Ontario by Koonce (1987). It contains two important changes, however. First, the ammocete submodel has been modified to allow for explicit treatment of the 49 lamprey producing streams in Lake Ontario. The rules for ammocete population dynamics are unchanged, but the spatial representation is quite different. Second, the simulation model was modified to update the stream database discussed above, after each simulated year. Since the database contains several dynamic variables (e.g., treatment date, ammocete densities), it has to be updated annually, according to both the population dynamics of lamprey and the implementation of the treatment schedule. Without this updating, the stream treatment strategy could not be implemented in subsequent years.

2.2.5 Expert System Report: Explanation Facility

When the simulation model has been executed for the desired number of years, control is passed back to the user, who has the option of repeating the sequence described above, viewing the results of the model runs just completed, or generating a report that describes all the choices made by the user in selecting the stream treatment strategy just implemented. Due to a lack of time, the explanation facility that exists within the prototype is quite primitive and not very useful. Although all the elements of the treatment strategy are documented, there is no explanation of how the chosen strategy resulted in the stream treatment schedule implemented in the simulation. Furthermore there is no explanation of how the model rules work. These are significant shortcomings of the prototype, but it is our view that with sufficient development work, a far superior explanation facility could easily be developed.

2.2.6 Results Presentation

Finally, the user will almost certainly want to see the results of implementing the strategy he has chosen. Each year the IMSL model saves the value of all of the management goal measures listed above, including the treatment schedule and control costs. At the end of the simulation, these results can be read into a LOTUS 1-2-3 worksheet template that has been set up to facilitate the generation of time series plots for each measure. Examples of these plots are presented in Section 3.

3.0 RESULTS

3.1 Application of the Prototype: Examples

In this section we present the results of four applications of the prototype. All four examples use the same stream selection criteria but each uses a different stream ranking algorithm. For stream selection criteria we selected a minimum three year treatment cycle and an ammocete density (ammocetes > 125 mm) of 0.05 m^{-2}. Thus only streams not treated within the past three years and with ammocete densities (> 125 mm) in excess of 0.05 m^{-2} were considered eligible for treatment.

3.1.1 Fixed Budget Algorithms

To select streams according to a maximum annual budget requires specification of costs of treating each stream. The stream database includes treatment area and chemical requirements per unit area treated. Assuming that chemical usage is proportional to total treatment costs, the expert system calculates the cost of treating each stream in the ranked candidate list. Streams are then selected in order from the ranked list until the budget is expended. If inclusion of a stream causes total costs to be greater than budget, it is omitted, and the expert system proceeds to examine lower ranking streams until the budget is spent. In the examples presented below, we assume an annual budget of $500,000.

The two methods of ranking candidate streams produce different annual stream treatment schedules (Figure 2). The cost/benefit ranking procedure produces quite regular fluctuations in area treated, with a maximum treatment of 40% and 45% of the total stream area. In contrast, the maximum benefit method apparently concentrates on fewer high cost streams and it treats no more than 20% of the total stream area in any year.

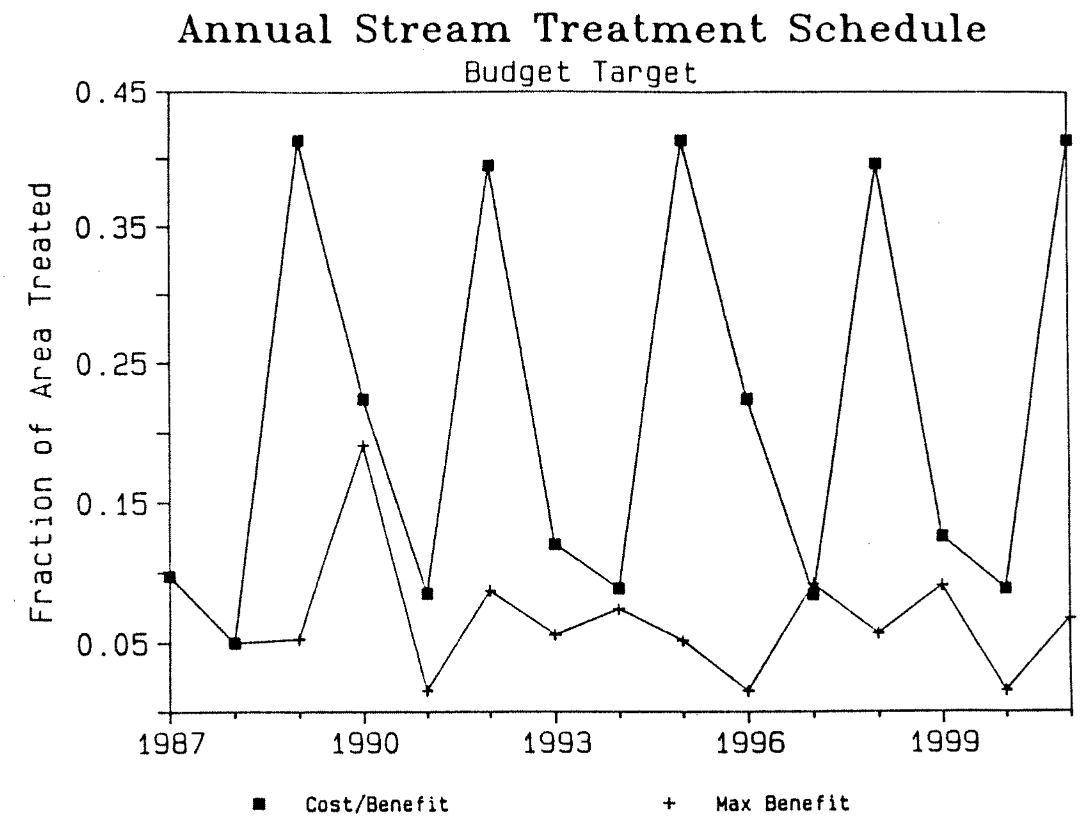

Figure 2: Fraction of total stream area treated each year in a simulation of fixed budget sea lamprey control for two methods of stream ranking (cost/benefit ratio and maximum benefit).

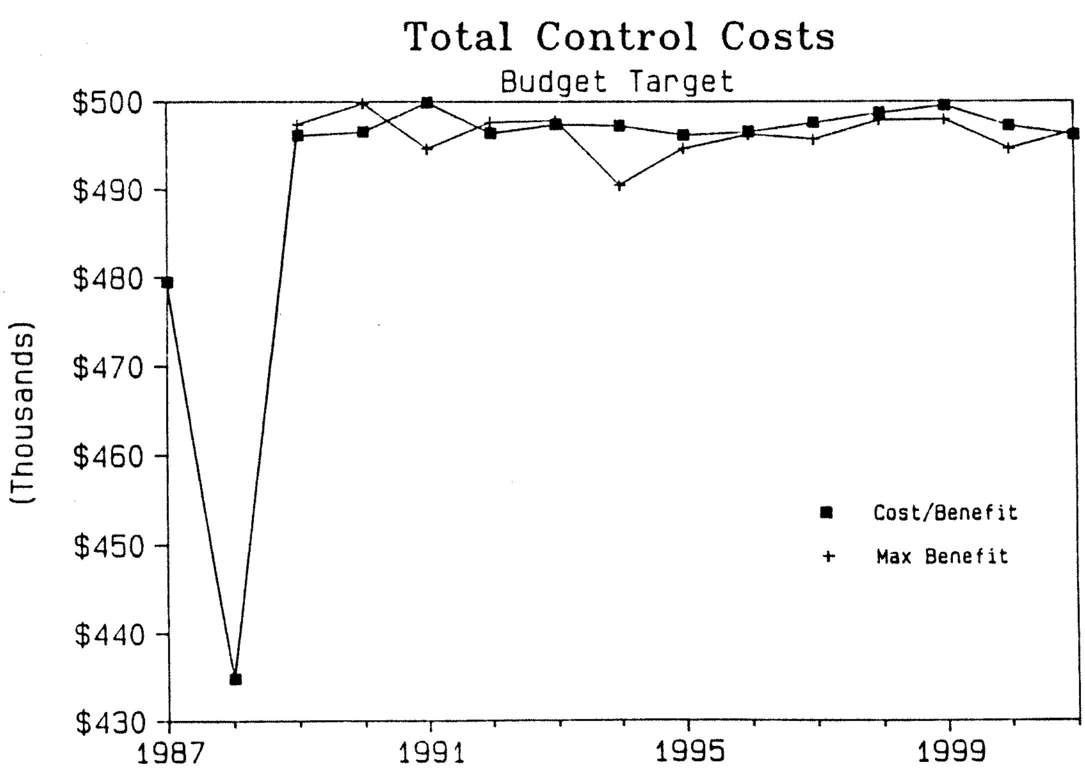

Figure 3: Annual control costs associated with a fixed budget control method for two types of stream rankings: cost/benefit ratio and maximum benefit.

As would be expected for a fixed budget method of stream selection, the annual control costs are nearly constant (Figure 3). Both methods of ranking streams produce almost identical schedules of annual cost.

Under the constraints of a $500,000 annual budget, spawning phase sea lamprey abundance increases for both ranking methods (Figure 4). The maximum benefit ranking method, however, allows a progressively higher residual population of adult sea lamprey. These results are counter-intuitive. The U.S. Fish and Wildlife Service control agents have recently evaluated various methods of selecting streams and they specifically cautioned against the use of a cost/benefit ranking method because it made major producers vulnerable to exclusion under a restricted budget. The opposite pattern seems to emerge in these simulations.

The rehabilitation target for lake trout in Lake Ontario is 0.5 to 1 million adult lake trout by the year 2000. Clearly both ranking methods meet this goal within a few years (Figure 5). In these simulations, however, fishing is assumed to be regulated by a strict quota and the level of recovery may not be nearly as great for more realistic fishery options. In any case, the recovery gained by 1998 does not appear sustainable under conditions of increasing abundance of adult sea lamprey (cf. Figure 3).

3.1.2 Target Ammocete Reduction Algorithm

Using a control target based on residual parasitic phase sea lamprey implies unlimited budget for sea lamprey control in any specific year. For this algorithm, the user must specify the residual target relative to the abundance of parasitic phase sea lamprey in Lake Ontario in 1986 (approximately 50,000 animals). The expert system estimates the residuals left after treating each stream and it continues to include streams for treatment from the ranked candidate list until the target residual level is attained. In

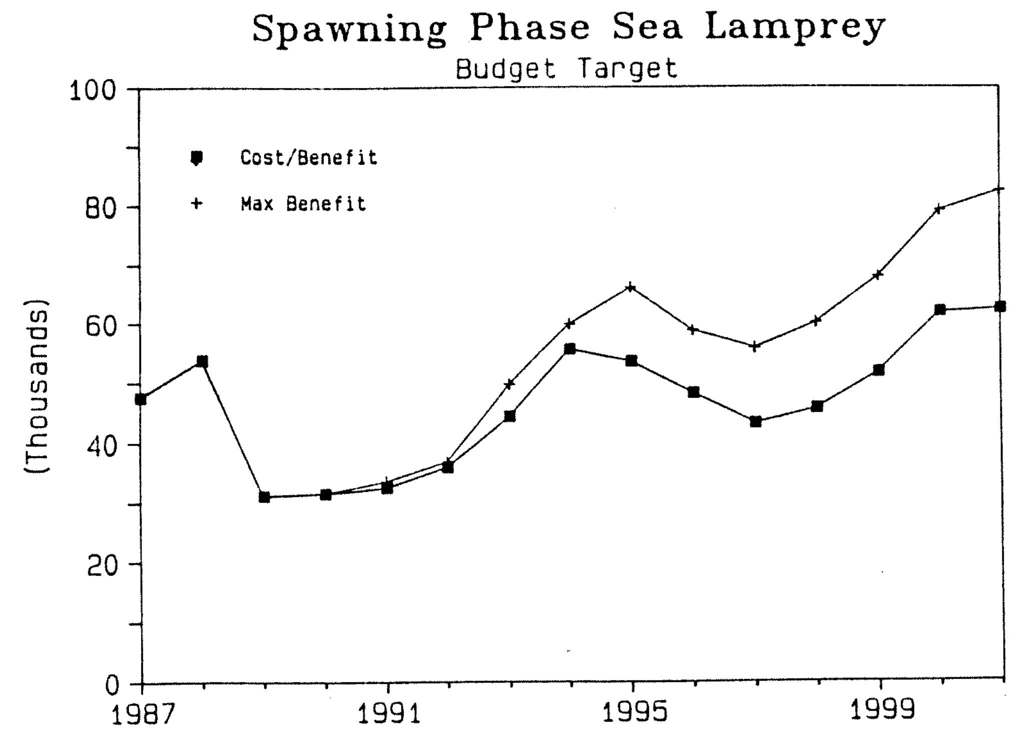

Figure 4: Predicted abundance dynamics of spawning phase sea lamprey for fixed budget control policies that rank streams either on a cost/benefit or maximum benefit basis.

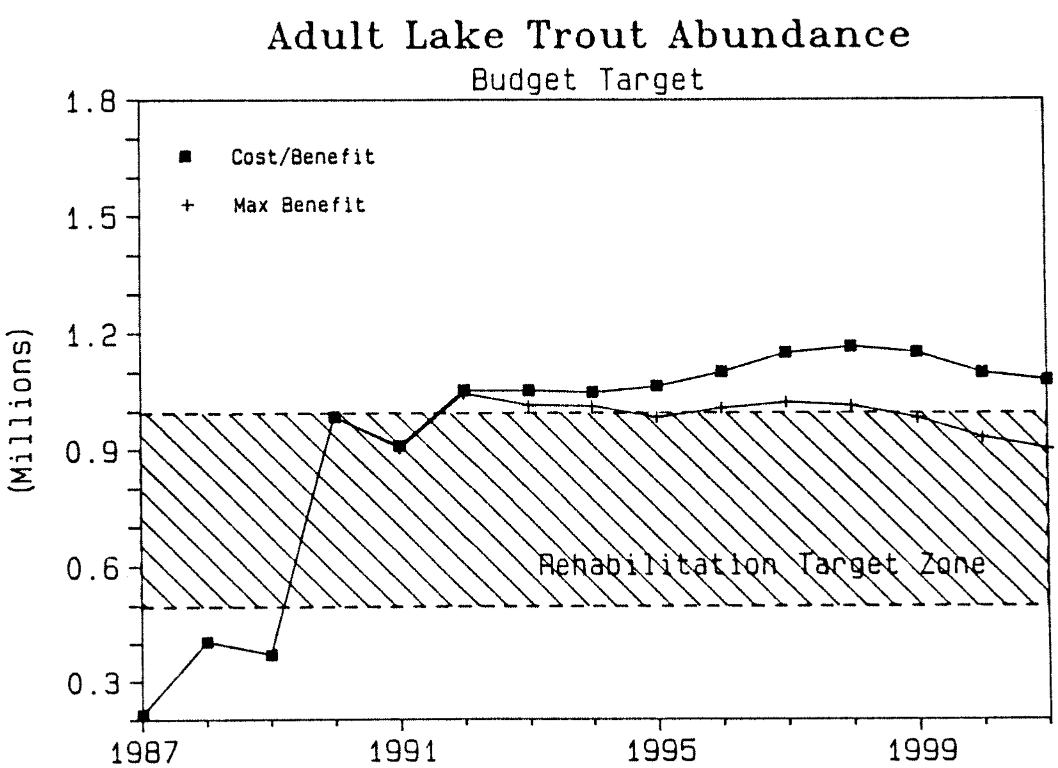

Figure 5: Predicted dynamics of adult lake trout for fixed budget control policies using cost/benefit ratio or maximum benefit as the basis for stream ranking. Fishing harvests are limited to an annual quota of 30.000 after 1990.

these examples, we found little difference between ranking methods (cost/benefit vs. maximum benefit), so we have focused instead on strategies oriented to reducing parasitic abundance in the future or maintaining the present levels.

The treatment schedules for parasitic phase targets (Figure 6) are quite different from those for a fixed budget (cf. Figure 2). To lower parasitic phase abundance seems to require total treatment in some years. Both schedules also seem to imply major treatment every fourth year.

The changes in annual control costs reflect the variation in treatment schedules (Figure 7). Lower lamprey targets would be more expensive (nearly $5,000,000 in some years) but the average annual costs are not as extreme. For lower lamprey, annual costs would average about $1,100,000 and the alternative would cost about $700,000 per year.

The spawning phase abundance patterns reflect the residual target levels (Figure 8). Treatment cycle variations will remain in the constant lamprey residual but the lower lamprey levels appear to be limited to a treatment effectiveness or survey detection threshold.

Sustained lake trout recovery (Figure 9) is a major difference in this method of stream selection. The implication is that far higher fishery yields will be possible using this method.

3.2 Limitations of the Prototype

In its current state, the prototype expert system has three major limitations that hinder its utility in an operational context. Two of these limitations were identified earlier: the stream database probably does not include all the information necessary to select among streams and the explanation facility is inadequate. The third limitation is that the system does not allow the user to influence the other management actions identified in the planning meeting (i.e., lake trout stocking and harvest controls). As with

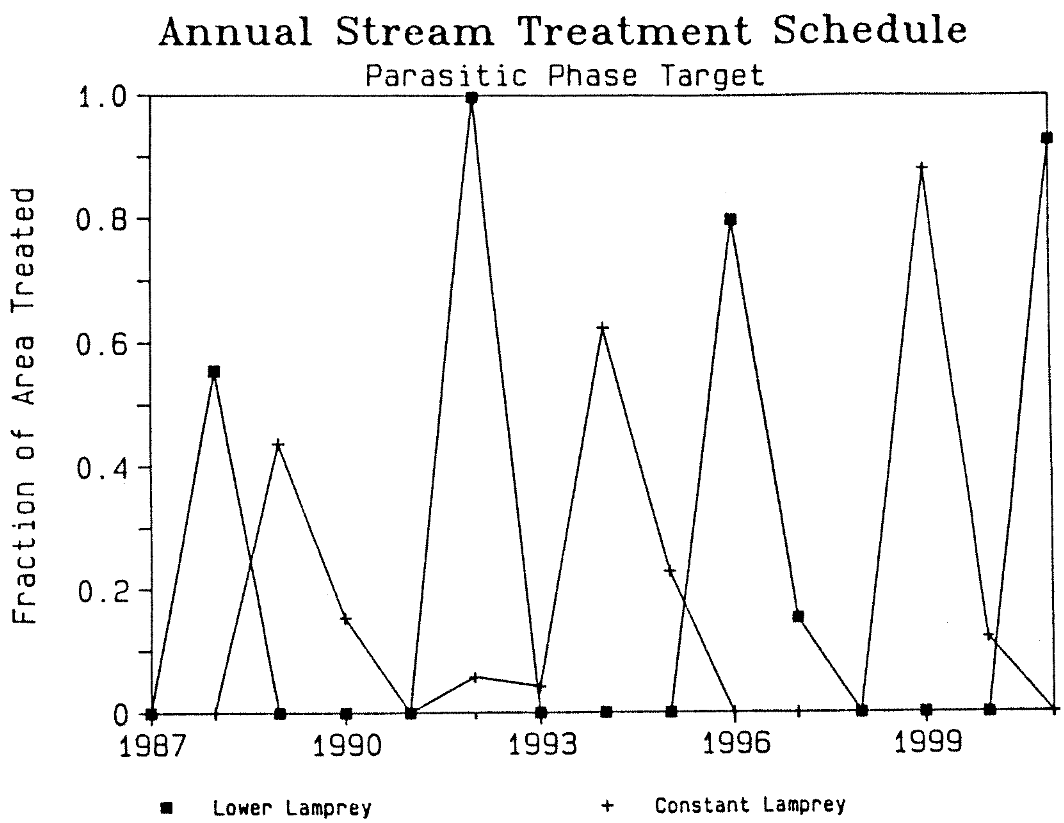

Figure 6: Fraction of total stream area treated each year in two simulations with differing target levels of sea lamprey control.

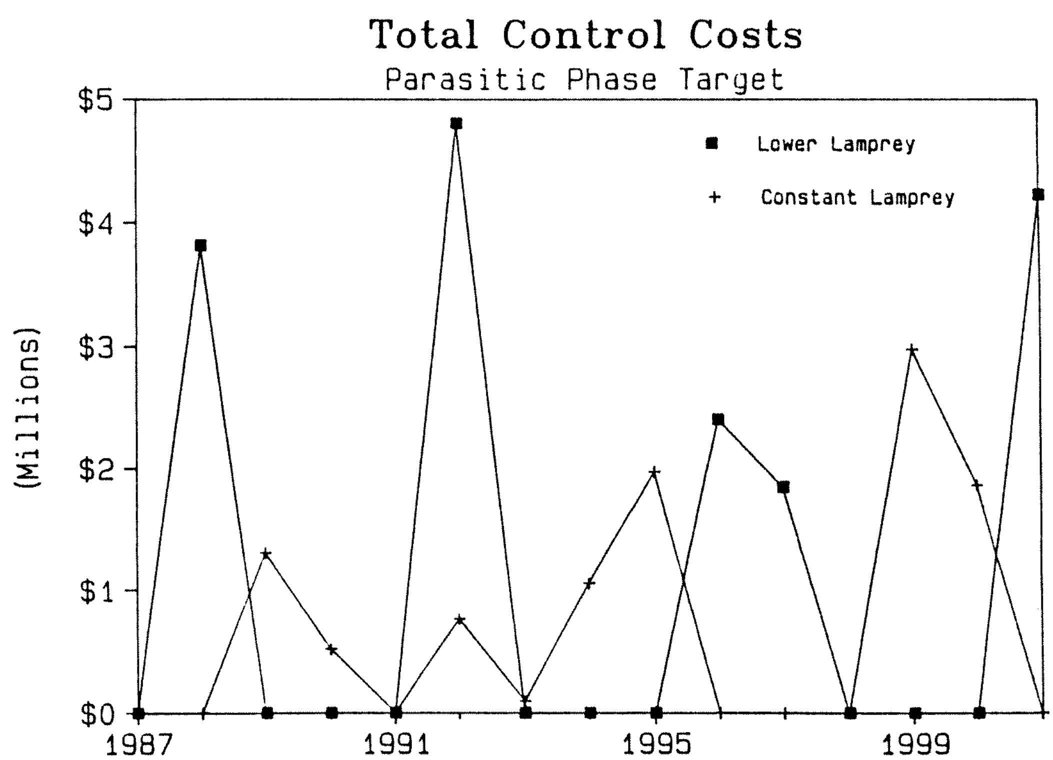

Figure 7: Annual control costs associated for the two simulations of target levels of sea lamprey control.

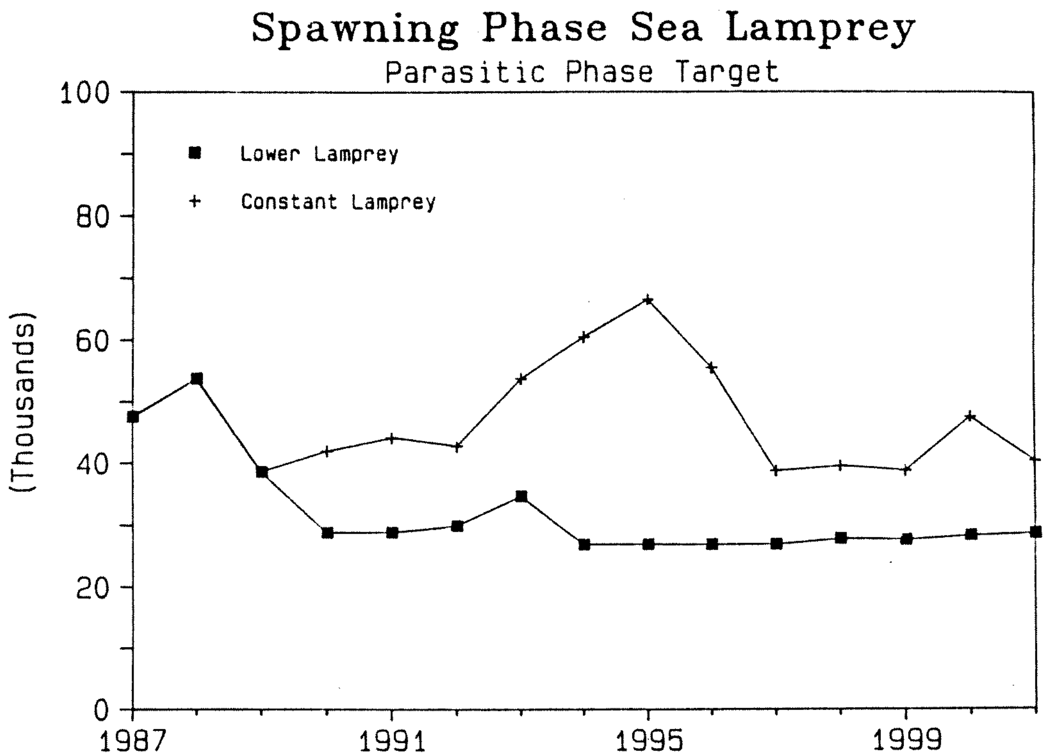

Figure 8: Predicted abundance dynamics of spawning phase sea lamprey for two simulations with differing target levels of sea lamprey control.

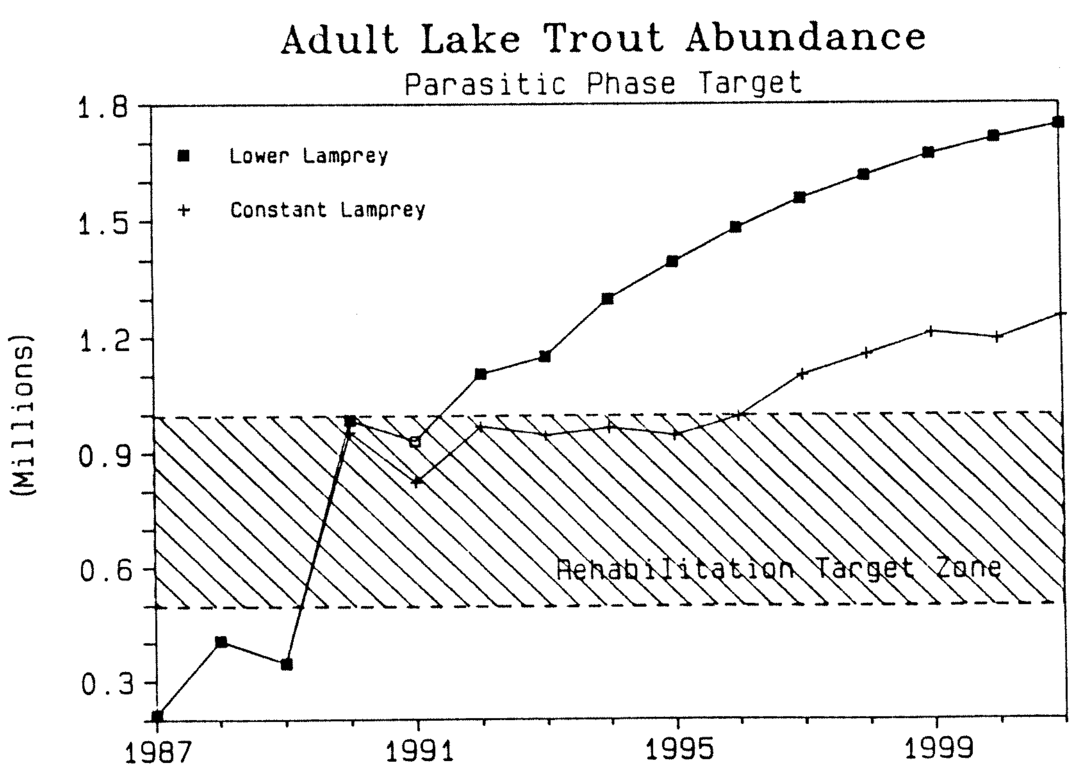

Figure 9: Predicted dynamics of adult lake trout for two simulations with differing target levels of sea lamprey control. Fishing harvests are limited to an annual quota of 30,000 after 1990.

the other two limitations, it is our contention that these actions could be incorporated into the prototype provided we were able to obtain sufficient guidance from fishery managers as to what the options for stocking and harvest controls might be.

3.3 Evaluation

At the project planning meeting in June, the steering committee identified seven possible criteria for evaluating the expert system (Table 4). Although we recognized at that time that most of the criteria in Table 4 would more appropriately be used in the final evaluation of an operational expert system, we still felt it would be instructive to judge at least the potential of the prototype to meet these criteria. Thus, at a second meeting held on September 11 in Detroit, the prototype was demonstrated to the steering committee, after which we discussed each of the seven evaluation criteria. In this section we summarize those discussions.

Comparison to Expert Judgement

The expert system did not choose, for treatment in the first year of the simulation, the same streams that were actually chosen in 1986 by the control agents. In this sense the prototype did not mimic the decision process used by the experts very well. As we pointed out earlier, however, the stream database probably does not contain all the information used by the control agents to select streams for treatment. Therefore it is hardly surprising that the prototype made different choices.

With further development, it would be useful to judge the performance of the system by supplying both the system and the current experts (i.e., the control agents) with the same information and comparing the treatment decisions made by each. If the two differed, it would be critical to learn why the differences arose, as this might point to either

Table 4: Evaluation criteria for the expert system identified at the June planning meeting.

1. Performance relative to expert judgement

2. Ability to explain logic used to reach conclusions: explanation facility

3. Technical credibility

4. Ease of use

5. Ease of adding new knowledge

6. Portability of software

7. Compatibility with other decision-making tools

flaws in the expert system's logical structure, or inconsistencies in the control agents' decision rules. Either way, the results would be highly informative.

Explanation Facility

We have already noted that the present explanation facility for the prototype is inadequate. Nevertheless it is clear from our limited work in this area that documenting the logic used by the expert system to reach conclusions is relatively easy, but time consuming. Providing an explanation facility for the IMSL model may be significantly more difficult.

Technical Credibility

In spite of the differences between the predicted and actual stream treatment schedules for 1986, we would suggest that as a prototype, the expert system has quite high credibility. Given the information that was available for development of the prototype, we feel the system reflects much of the state-of-the-art in the management of sea lamprey and the effects of sea lamprey control on lake trout. At the technical level, the system's greatest weakness presently is the inadequate stream inventory database.

Ease of Use

In its present state, the prototype is quite easy to use, even by those unfamiliar with the use of microcomputers. The management of the interaction between the user and the expert system could be substantially improved, however, and should be if an operational system is to be developed. Although a carefully prepared user's guide would be essential, the ease of use of an expert system such as this is one of its most obvious assets.

Ease of Adding New Knowledge

The first stage of the expert system is explicitly designed to facilitate the addition of new knowledge that fits into the stream database structure. In this respect it is obvious that new knowledge is easy to add. New information that changes the logical structure of the problem (e.g., an entirely new criterion for stream selection), poses a much more substantial challenge. Such changes are not impossible to make, but would likely take the same amount of programming effort as a similar structural change to a simulation model. Ideally, careful attention to possibilities during the system development stage should minimize this problem.

Portability of Software

It is difficult to say at the present time whether portability will be a problem. The existing prototype uses INSIGHT II as the expert system shell, Microsoft FORTRAN v4.0 as the programming language for the model and some of the expert system subroutines, and LOTUS 1-2-3 for presenting results. A system that supports these packages could run the expert system. It is unlikely, however, that this combination would be the best choice for further development. If portability is an important issue for the Commission, it should be confronted before further development occurs. At this point we would tentatively recommend that the operational language be written entirely in one widely used language, such as C.

Compatibility with Other Decision Tools

Briefly, it is our view that this expert system is highly compatible with other tools either in use or under development. As with all such tools, the expert system should never be thought of as doing the job of a manager or decision maker; rather it provides a rational basis of support for that manager's decisions. These issues are discussed further in the next section.

4.0 CONCLUSIONS AND RECOMMENDATIONS

The overall objective of this study was to develop a prototype expert system that would serve to illustrate the potential use (or lack thereof) of such tools for the Great Lakes Fisheries Commission. In our view this objective has been achieved and the example applications of the prototype presented in Section 3 clearly demonstrate the potential value of this type of tool to activities such as IMSL. In this concluding section we first offer some additional thoughts on the possible uses of this tool given that resources are committed to its further development. Then we present four specific recommendations for further development based on discussions with the project steering committee at the September 10 evaluation meeting.

By its basic character, a properly developed expert system should make the same decisions as an expert. The application of an expert system to integrated management of sea lamprey, however, has the potential to make the rational basis of stream selection more explicit. Communication is an essential component of the control agents' job and the expert system can help the control agents present the rationale for their decisions and budget needs in a more objective manner. The expert system can be used by others in the Great Lakes Fishery Commission to explore the trade-offs implicit in the decisions and policies of the sea lamprey control program.

An expert system, like this demonstration, helps the control agents screen candidate streams during routine scheduling. The expert system provides an explicit framework for integrating historical data and current survey information. Under ideal conditions, the selection of streams would be a simple matter once basic decisions are made concerning criteria for treatment. In most cases, however, unanticipated factors (dam washout, extreme risk of mortality to sensitive non-target species, etc.) require modification of the treatment schedule. The expert system accommodates such modification and provides a way to explore the consequences of alternative schedules.

Ultimately, IMSL will involve a host of trade-off decisions. Stream selection for chemical treatment is only one component of these decisions. Nevertheless, the expansion of this expert system to include a wider range of control and fishery management options would provide a framework for joint decision-making and evaluation of alternative management decisions.

4.1 Recommendations

The project steering committee agreed that the prototype seemed to have sufficient merit to warrant some further expansion. Any expansion, however, must be carefully coordinated with the developing efforts to implement IMSL within the Great Lakes Fishery Commission. In this context, we offer the following recommendations.

The stream treatment schedule recommended by the expert system deviated significantly from current practices, according to the Canadian control agent on the steering committee. Much of the deviation is caused by the lack of realistic survey data on ammocete densities in the stream inventory database. Lacking explicit survey data, the database contains estimated ammocete density and age structure based on its history of treatment and assumed recolonization patterns for small, medium, and large streams. These "estimates" were arbitrary starting points to demonstrate the expert system but they could be updated with better use of existing survey information. We recommend, therefore, that the development of a stream inventory system (including basic physical and chemical characteristics of the streams and estimated ammocete density and habitat area) be made a high priority task for the control agents. The improvement of the stream inventory database will be required before major enhancement of the expert system is possible.

This prototype expert system is a possible component of a more general decision support system. A decision support system for IMSL should provide an aid to explore the consequences of various trade-off decisions involved in the rehabilitation of Great Lakes fisheries and sea lamprey control. By providing a framework for synthesizing monitoring information, a decision support system is a vehicle for evaluating policy options, justifying budgets, and allocating resources. Some aspects of a decision support system may benefit from the use of expert systems. Expert systems should be limited to those aspects of the problem for which decision rules are clear enough to warrant logical analysis. It is very unlikely that the entire IMSL program could be included in an expert system, but our experience with this prototype convinces us that a more comprehensive decision support system is feasible. We recommend, therefore, that the Great Lakes Fishery Commission further explore the incorporation of expert systems into IMSL decision support systems currently under development. We do not see any technical impediments to simultaneous considerations of individual stream treatment schedules with other sea lamprey control options and fishery management initiatives.

The prototype expert system focused on stream treatment schedules for a single lake. The control agents, however, are responsible for coordinating treatment in the entire Great Lakes Basin. The logical structure of basin wide decisions, as well as the fundamental need to develop more explicit guidelines for allocation of control resources, create a different set of issues with which an expert system must deal. We recommend, therefore, that future expansions of the expert system deal explicitly with the problems arising from basin wide concerns. Building technically more credible expert systems or decision aids while ignoring the dominant problems of allocation of control resources, target levels of control, and budget justification will inevitably result in major delays in applying the decision aids.

Although promising, major uncertainties about the ammocete phase of the sea lamprey life cycle will continue to compromise the development of decision support systems or expert systems to aid IMSL. We recommend, therefore, that Lake Ontario continue to be a test case for improving the technical credibility of the models upon which both this expert system and developing decision support systems are based. Many of the findings of additional work will apply to the other lakes as well and can be incorporated into basin wide considerations without much delay. This effort is also required to determine the appropriate level of resolution for assessment and monitoring databases.

5.0 LITERATURE CITED

Koonce, J.F. 1987. Application of models of lake trout/sea lamprey interaction to the implementation of integrated pest management of sea lamprey in Lake Ontario. Final report submitted to Great Lakes Fishery Commission. 27 pp.

Spangler, G.R. and L.D. Jacobson (eds.). 1985. A Workshop Concerning the Application of Integrated Pest Management (IPM) to Sea Lamprey Control in the Great Lakes. Great Lakes Fishery Commission, Special Publication No. 85-2. 97pp.