ELSEVIER

Ecological Modelling 75/76 (1994) 617–628

ECOLOGICAL MODELLING

Using machine learning techniques in the construction of models

I. Introduction

B. Kompare [a,*], I. Bratko [b], F. Steinman [a], S. Džeroski [c]

[a] *Department of Hydroengineering, University of Ljubljana, Hajdrihova 28, P.O. Box 579, SLO-61101 Ljubljana, Slovenia*
[b] *Laboratory for Artificial Intelligence, University of Ljubljana, Ljubljana, Slovenia*
[c] *Jožef Stefan Institute, Ljubljana, Slovenia*

Abstract

A research was initiated in automated modelling of the ecosystem using deep knowledge and machine learning techniques. The goal of the research is to show that using advanced artificial intelligence (AI) techniques, measurements, and some general basic knowledge about the ecosystem suffice to automatically generate better models and in less time than is the case by traditional construction of models. Namely, advanced techniques of AI are able to identify and model a system that we do not understand yet. The methodology of the approach is presented and illustrated by an example of a successfully constructed model in the field of medicine. Subsequently some encouraging results were obtained in the field of ecology and related sciences.

Key words: Artificial intelligence; Machine learning; Model building

1. Introduction

The advent of computers and in particular personal computers has brought the possibilities of mathematical modelling of ecosystems to every modeller's desk. From the first statistical (black-box) models almost everybody proceeded to deterministic, conceptual or physically based (i.e. transparent, or white-box) models,

* Corresponding author.
An earlier version of this paper (by Kompare, Steinman and Bratko) was presented at ISEM's 8th International Conference on the State-of-the-Art in Ecological Modelling.

SSDI 0304-3800(94)00024-C

since the latter give more insight into mechanisms. Due to this reason, the deterministic modelling is most widely accepted and used, although it may not be the best way of describing and/or modelling a real system, taking ecosystem as a special, complex case. One difficulty is that making models more detailed, one needs more and more data to calibrate and validate them to be useful. This topic of complexity and expressiveness of the model was discussed many times (Jørgensen, 1986). Another difficulty in deterministic modelling is that the processes themselves, along with the measured data, are subject to different perturbations which can be roughly classified as white noise when random (coming from the outside, i.e. stochastic), or as pink noise when caused by an unaccounted for internal reaction of the system (Beck, 1978, 1983; Bossel, 1992). A sound basis to determine the structure (and appropriate complexity) of a deterministic model is to identify it with the systems analysis tools (Beck, 1978, 1983; Bossel, 1992; Young, 1992). It turns out that for the construction of complex dynamic models, basic (deep) knowledge is only a necessary, but not a sufficient condition. Complex systems are not a simple sum of mechanistic subsystems which can be deducted from isolated measurements, but are rather holistically inter-bounded (one would say synergistic). That is the reason that intensive and complex statistical analysis, together with other methods of systems analysis, are also needed. Such an approach can then identify and depict the very important, but too often overlooked, indirect effects among subcomponents of the system (Patten, 1982; Patten et al., 1990). The importance of this phenomenon is perfectly depicted as the butterfly effect.

Even the easiest way of constructing deterministic models out of only deep knowledge is for a majority of ecologists a rather difficult task. It is hard to think and express in strict mathematical formalisms, and even more difficult to code this mathematics in a computer language. And if they pass these tests, there are still traps of numerical modelling, lack and inconsistency of data, difficulties to determine the exact modelling parameters, etc. All of us would appreciate a computer program being capable of understanding our common descriptions of ecosystems and transforming them into a suitable machine code. Some attempts were already made to simplify the construction of ecological models by constructing a shell, or a special language, in which description of the model and parameters are easily input, i.e. STELLA (High Performance Systems Inc., USA) on Macintosh, BSIM (B. Silvert, Canada) and SYSL (E^2 Consulting, USA) on VAX and PC, SIMSAB (ECOLEX, Russia) and LAKE (Jørgensen, 1992) on PC, etc. Despite these tools, there is still not much of automatism in constructing the models. Pioneering approaches in the field of automated construction (machine modelling the modelling processes) were done by Muetzelfeldt et al. (1989), who are developing programs to help in the construction of simulation models. Although this approach represents a giant step in simplifying the process of model construction it is still rather exhaustive for the ecologist by asking him/her quite a bulk of questions, to which sometimes no decisive answer can be given, due to the lack of information. On the other hand there might be a lot of data which are difficult to digest by ordinary ecological experts (or any human), but could possibly be analysed by some intelligent software, which may (we believe should) find some meaning in them; of course, taking into account that at least the crucial processes are adequately measured. If such a software does not make only various statistics, but also combines known basic facts (background, deep knowledge) about the ecosystem which one is trying to model, then the constructed model should not be just a simple black-box model, but more or less transparent, i.e. grey-box. With the term grey-box we are trying to depict that the model is relying on basic principles, although some relations are not completely transparent (disclosed) and are merely given as the result of statistics. This kind of model construction and modelling is the very subject of our research.

To our knowledge very few attempts have been done in this direction, i.e. Šterbáček et al. (1990a,b), Guariso et al. (1992), Young (1992) and Dejak et al. (1993). Šterbáček et al. are developing an expert system (ES) shell in which they use ready-made knowledge base (KB) on which they superimpose heuristic reasoning. In the KB there are already made and tested crisp mathematical submodels. With the heuristic approach the authors try to grasp the fuzziness of the real system, i.e. synergistic combination of the crispy submodels. Young is working in the intermediate space between a mechanistic and a statistic (data-based) approach. He exploits the availability of time-series data in statistical terms but with the attempt to automatically produce models which have some sensible physical interpretation. He accepts such a model only if, in addition to explaining the data well, it also gives a relevant description of the physical processes of the studied system. Guariso and coworkers have developed a QUALSIM shell, which helps the expert to identify the structure of a dynamic system. The package allows qualitative modelling, so the final definition of all the parameters remains still open. But, as stressed by the authors, the main scope of QUALSIM is to support the model-building process by enforcing the user's intuition. The process is highly interactive, as it needs user intervention at each point of ambiguity, i.e. usually when qualitative derivative (decreasing, steady, increasing) changes its value. The algorithm used is heuristic and uses a fuzzy qualitative-numeric conversion that avoids the definition of a qualitative algebra. One very practical result of this modelling tool is that it is quick and it is easy to try and evaluate different structures with small structural perturbations. This is more useful than time-consuming fine tuning of an inadequate numeric program. Dejak and coauthors have tried the approach to identify the structure of the model by use of Fourier series. They have also determined optimal complexity of the model by the use of minimum negentropy criterion. Although qualitatively the agreement between measured and simulated variables is good, there is still a lot of work needed to achieve acceptable quantitative fit. Although not in the same way as we defined our scope of developing (ecological) models, Salski (1993) is exploring the possibilities to use fuzzy-knowledge-based models in ecology.

It is important to stress that major improvements in this approach can be only made with the help of the methods of artificial intelligence (AI). We here assume the basic knowledge about first-generation expert systems (ES). If this is not the case, one should refer to the description of ES elsewhere – as a suggestion we mention a few, i.e., Frenzel (1987), Turban and Watkins (1988), Bratko (1990). The following are relevant in particular with respect to ecological modelling: Davis et al. (1989), Reinhardt et al. (1989), Ritchie (1989), Starfield et al. (1989) and Šterbáček et al. (1990a,b). Expert systems of the first generation, which are most known to majority of us, usually rely on the use of surface, or shallow (i.e. heuristic) knowledge. Besides this the knowledge acquisition by interviewing and querying experts was a rather exhausting task for an unskilled modeller. This is the main reason that the great fame and expectations about ES have slowly turned into quiet disillusion. To the contrast, the second generation of ES, which is now in progress, is usually characterised by two additional features: deep knowledge and machine learning. Deep knowledge enables an expert system to reason from basic (physical, chemical, etc.) principles, which improves the system's robustness, explanation capability, and verifiability. Machine learning to great extent reduces the need to query the expert in the way that machine extracts the knowledge from the given examples (data). Both features of the second-generation ES make an ES capable of identifying and modelling a real-world system that we do not understand yet. In relation to deep knowledge, and opposed to the traditional quantitative modelling (i.e. everything expressed by numbers), qualitative (i.e. descriptive) modelling is being recognised as increasingly important, especially in the fields where only qualitative information is known (e.g. common sense reasoning, fuzzy and noisy data). In this paper we will focus on the description of the implementation of some methods of this second-generation ES with a particular emphasis on deep knowledge and machine learning.

2. The second-generation expert systems – the KARDIO approach

This article is aimed to be an introductory paper to the various methods and principles of AI which can find usage in ecological modelling. Recently we have obtained some very interesting results with different approaches and methods of the second-generation ES. These results will be briefly illustrated later in this paper, and presented more in detail in subsequent papers. In this paper we focus on the special knowledge acquisition cycle KARDIO (named after Bratko et al., 1989), which is becoming a standard constructional element of the methods of the second-generation ES. The underlying principles are shown in the following.

We start the discussion by stating the differences between *surface* and *deep* knowledge. Surface, or shallow knowledge directly states the relationship between problem specification and problem solution without referring to the underlying principles. Surface knowledge is typically used in solving problems efficiently, although without any reference to, or understanding of, the underlying causal relations on which the solution is based. Another name for surface-level knowledge is thus also "operational knowledge". An example of shallow knowledge quantitative models are statistical and/or empirical black-box models. Deep knowledge relies on basic (also so called *first principles*) of the specific domain and on needed basic knowledge from other disciplines, primarily mathematics, physics, chemistry, biology, etc. Models dealing with such kind of knowledge are called white-box, conceptual, physically based, etc.

The distinction between shallow and deep knowledge can be illustrated by human problem solving. When solving a problem from their domain, human experts usually do not have to think hard about the solution – they already *know* the right (or sufficiently close) answer. They simply retrieve the answer from their "operational knowledge base". Only when faced with an unusual problem, they can not pick up the solution from their experience, but have to synthesise or derive it by reasoning from the "first principles". Applied to AI and ES, this is the distinguishing feature of the second-generation ES.

It is obvious that not all forms of knowledge are equally suitable for all tasks that ES are supposed to perform. Regarding the direction of inferencing, some representations are better suited for answering prediction-type questions that involve reasoning from causes to consequences (forward chaining); whereas other forms better suit diagnostic questions that involve reasoning from consequences (known manifestations) back to the causes (diagnoses), which is called backward chaining. And regarding the quality of the knowledge, the task of explaining may impose on ES to reference the underlying mechanisms, i.e. deep knowledge. On the other hand, for simple-minded and efficient diagnosis surface knowledge might be sufficient or even preferable to deep one, since typically the use of surface knowledge is operationally more efficient.

It is therefore useful if not obligatory to use different representations of the same knowledge for different types of tasks. Once we have some representation, it is most convenient to automatically transform this representation into the one which is best suited for the particular task. In the KARDIO study (Bratko et al., 1989) the aim was to develop a KB capable of diagnosing cardiac arrhythmias from ECG traces. In that study the following transformations were carried out (see Table 1 and Fig. 1).

In transformation 4, the system in fact automatically constructs a theory from data. The first three transformations were completely automatic whereas the fourth one was semi-automatic and required human interaction. This interaction was necessary because normally surface knowledge does not sufficiently constrain the deep model. Namely, various deep models can explain the given surface behaviour, so additional constraints on the deep model are needed.

The first transformation was done with exhaustive simulation based on deep model by deduction. In this step, intensionally represented deep knowledge

Table 1
Transformations of knowledge representations

Transformation No:	1. deep model	→ surface knowledge
	2. surface knowledge	→ compressed diagnostic rules
	3. surface knowledge	→ compressed prediction rules
	4. examples of surface facts	→ deep model

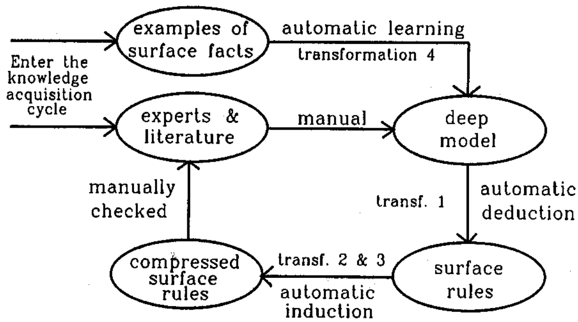

Fig. 1. The KARDIO knowledge acquisition cycle (rearranged from Bratko et al., 1989).

(mathematical conceptual model) is transformed into extensive surface representation of the same information. The deep model can be viewed as a set of logical axioms from which the surface knowledge logically follows as a set of theorems. These theorems can be presented as a list of simple rules. Such a list is evidently much bulkier than the original deep knowledge base. However, the information required by the diagnostic or prediction-type questions is stated explicitly in the table and can be simply retrieved. Such a representation of knowledge is convenient for machine use but can be too complex (due to its size) for humans to study and to understand as a whole. Step 1 thus has equivalents in ecological modelling when having a model and constructing all possible outputs to given possible inputs. If the model is not very simple, the vast output cannot be (easily) digested by humans, unless compacted in some way. The way of compacting the shallow knowledge (output of the model) is the matter of next step of the KARDIO cycle. In order to get more compact but still shallow representation, inductive learning was used to compress this data base into compressed prediction and diagnostic rules (transformations 2 and 3).

In transformation 4 individual facts from the surface knowledge are used as examples for a learning program. These surface facts input to the learning program can be also obtained experimentally as results of measurements. The learning program then formulates compact rules defining a particular concept or relation, e.g. the behaviour of a heart's component in the case of the KARDIO project. The inductive program generalises and obtains rules that are more general than the examples. This means that the procedure is not necessarily truth preserving. However, if the program uses all the rules from surface knowledge-base as learning examples, and thus covers the whole problem space, and if the program has the property of perfect reproduction of the learning set, then the learned rules are a compressed representation of and equivalent to the original surface knowledge. The compressed rules obtained in transformations 2 and 3 are compact and easy to understand. The two mentioned compressed representations are thus most suitable for human experts to study and to compare with the existing human codifications of the same knowledge, which is invaluable in knowledge validation. Differences

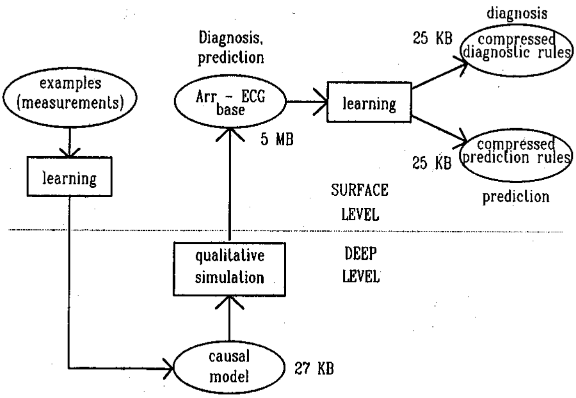

Fig. 2. Deep and surface levels of cardiological knowledge and transformations between these representations (rearranged from Bratko et al., 1989).

between the original human codification and the compressed representation can be due to either: (1) an error in compressed representation, or (2) a slip in the human codification. In the first case the error is in the deep model since both the deductive and inductive derivation steps are truth preserving due to the implemented mechanisms. In this case the knowledge acquisition cycle provides a feedback loop for verifying and correcting the deep model (see Fig. 1). In the second case the knowledge acquisition cycle will expose blemishes in the existing original expert formulations (e.g. in the literature) and may thus help to improve those.

There were many practical achievements of the KARDIO project. We list only those which have most relevance to ecological engineering:

(1) The bulky surface knowledge about diagnostic or prediction-type questions was compressed into suitable and easy to analyse and understand compact rules. Speaking in terms of computer memory, 5 Mbytes needed for the representation of surface KB were compressed 200 times to only 25 Kbytes of diagnostic rules, and 25 Kbytes of prediction rules. See Fig. 2.

(2) The learned system produces the same or better diagnostic answers than the domain experts. This finding was shown to be valid also in other cases of implementation of advanced techniques of AI (see Bratko, 1993).

(3) The usefulness of the developed compressed KB is best demonstrated by the fact that part of this KB is implemented in an artificial pace maker manufactured by a medical instrumentation company.

The KARDIO project has shown that the described knowledge acquisition cycle, involving qualitative modelling and machine learning technology, can be used to construct knowledge bases whose complexity is far beyond the capability of traditional, dialogue-based techniques for knowledge acquisition. Authors of the method (Bratko et al., 1989) believe that this approach should become a standard method in the development of practical expert systems.

3. Possibilities to use the KARDIO approach in ecological modelling

In the case of KARDIO, a mathematical description of the heart viewed as a mechanical device with an electrical control system was constructed. This represented a deep model, i.e. deep knowledge. From this deep model the surface knowledge was produced via exhaustive simulations of different operational conditions of the model. The same approach is also often used in ecology – we first build a model and then play with it to understand the complex relations in the ecosystem. What is new in the approach of KARDIO is:
1. the closed knowledge acquisition cycle, and
2. the compression of surface knowledge/rule base.

ad 1. Besides the evident salient capability to help in model validation, as described in previous section, the closed knowledge acquisition cycle means indeed that all the forms of information in this cycle are equivalent. That is, each step in this cycle encompasses *all* the relevant information, but in a *different* way. This essentially means that it is in principle irrelevant at which step we begin the cycle. Instead of beginning with a deep knowledge model, one can begin with surface rules.

ad 2. The compression of the surface knowledge/rule base is of great importance from a computational point of view. But its real advantage is that compressed rules are easier to comprehend and understand by humans. And in the case when one starts the KARDIO cycle from the surface knowledge, compressed rules are a valuable source of information to construct the deep model.

The idea to use the KARDIO approach in ecological modelling is to start the KARDIO cycle from the point of surface rule base, or better, from examples of surface facts.

The real problem by this approach is to construct the mentioned surface rule base. It would be impractical to use a second-generation ES and perform simple surface knowledge acquisition by querying as in the first-generation ES. No, what we would like to do is to let the ES itself to construct a needed model from the data at our disposition (measurements, descriptive knowledge, etc.) and from the basic underlying principles of the problem's domain. This approach is called automatic or machine learning. The basic approach and methods were first developed and used in the field of pattern recognition and clustering of data. The first methods were merely more sophisticated statistical methods, and rule bases obtained in this way were difficult to understand and analyse by humans. So new techniques have emerged, where the result of learning can be expressed as a list of rules, formulas, theories, or as a description of the concept in a formalism, easy to understand by humans (e.g. graphically presented decision trees). In this way the modeller can see the relations, the underlying principles, logic of reasoning, etc. (Lavrač et al., 1989).

Some successful systems for inductive learning from examples have been developed at the University of Ljubljana and at the Jožef Stefan Institute in Ljubljana. The first one is known under the name ASSISTANT PROFESSIONAL (Cestnik et al., 1987). This ES can handle continuous and discrete data and can take into account uncertain, incomplete, and noisy data. The system can detect and eliminate inconsistent and exceptional cases from the example set and thus improve the structure and accuracy of the generated (learned) knowledge. Another learning system is RETIS (Karalič, 1992), which seems to be even more applicable in ecological modelling since it does not need to first classify the continuous data. The output of both programs are graphically depicted decision trees. A feature of both methods is that they can handle fuzziness to some degree, thanks to their ability to operate with noisy data. The third program called LAGRANGE (Džeroski and Todorovski, 1993) can identify system structure (i.e. algebraic and ordinary differential equations) and values of governing parameters, given only sufficiently long time series of measured data. There is danger that the learning system gives plausible theories which correctly describe and reproduce learning examples, but which are hard to understand and justify by the general principles applicable to the domain of modelling. This question first arose by complex models, i.e. when checking the results of RETIS on the data for the Lagoon of Venice.

Regression trees induced by RETIS from the data for the Lagoon of Venice were checked by two experts (Prof. G. Bendoricchio and Prof. S.E. Jørgensen). Both experts agreed that the knowledge generated by the program qualitatively fits to a very high degree their experience and general ecological knowledge of a well-trained (experienced) person for the specific case. Fig. 3 shows such an induced regression tree.

In the case of application of LAGRANGE, we still have problems generating adequate differential and algebraic equations describing the phenomena in the lagoon. First, too few forcing functions and state variables were measured, and second, some measurements were inadequate (measuring DO at noon just below surface), missing (tidal currents, wind, etc.), or insufficiently accurate (estimated errors of the order of 50% or more). In such conditions proper automatic system identification with the generation of differential equations is hardly possible. Still, test results on synthetic data show that the approach is promising – i.e. LAGRANGE discovers Monod equation, population dynamics, chemical reactions, mechanical principles etc. (Džeroski and Todorovski, 1993).

In ecology, another difficulty is often experienced; i.e. we do not have enough learning examples. Just imagine that we would like to construct a model of lake eutrophication. We can have years and years of measurements but all of them describing the same trophic state. From such a set of learning examples a learning system will not be capable to construct a model of behaviour of the lake in changed trophic conditions (e.g. change of structure of the system after the reduction of nutrients input). This lack of data is the second factor that has been impeding our research. But on the other hand, the same procedure implemented on a set of examples of different trophic states, or different events of special concern, can give good and useful results. We believe this particular approach should be used to find some sound relationships and mechanisms which govern appearances of jelly-fish or algal blooms in the Adriatic Sea (mare sporco). Hoping that only classical biological research work will give successful explanation is too optimistic, whereas the belief that so complex phenomena can not be modelled is

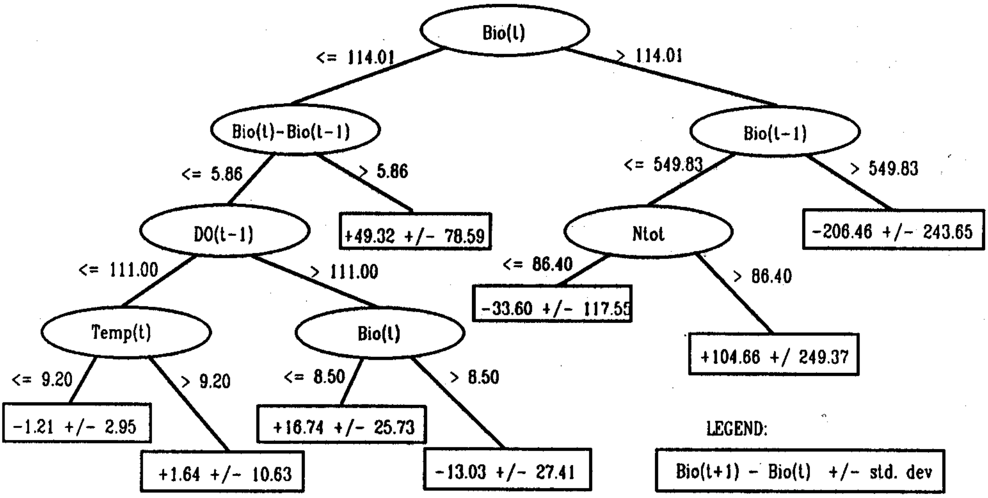

Fig. 3. Example of a regression tree induced by RETIS from the data for the Lagoon of Venice. The tree predicts the value of the algal biomass ("class") as a function of other parameters ("attributes"). Nodes of the tree (ellipses) correspond to the attributes, while leaves (rectangles) represent the value of the predicted class. In our case we wanted to predict changes in biomass in the following week ($\mathrm{Bio}(t+1)-\mathrm{Bio}(t)$), based on the measured data from this week (t) and the previous week ($t-1$). Measurements comprised five attributes: (1) algal biomass (Bio), expressed in g/m^2 of dry weight, (2) temperature of water (Temp) in °C, (3) dissolved oxygen (DO) in % of saturation, (4) total nitrogen (Ntot) in ppm, and (5) phosphorus (P) in ppm. The hierarchy, the presence and the values in the tree clearly depict the most important parameters, conditions when they are crucial, and their respective values; e.g. as phosphorus is not present in the tree but nitrogen is, we can easily conclude that phosphorus is neither limiting nor interesting nutrient for algal growth (at present conditions, of course). It means that it is always in excess, which was also confirmed by experts (Prof. Bendoricchio).

outdated and too pessimistic. We believe the answer to the above question is just at the grip of our hands – with the help of inductive learning techniques.

4. Conclusions

We have presented a methodology for computer-aided development of (ecological) models. The method is promising as indicated by the results from automated modelling of the Lagoon of Venice. We are also trying to implement this method in modelling two lakes in Slovenia, the Lake of Bled and the Lake of Bohinj. For the lakes there are no big changes in forcing functions (in reality and measured), therefore a predictive model for simulating response on greater changes in forcing functions perhaps will not be possible to develop with the described techniques. Experiments to model seasonal dynamics (algal blooms) are underway. According to our experience we strongly believe that the described approach will show its benefits compared to classical construction of (ecological) models regarding the ease of construction and calibration, and regarding the choice of adequate structure of the system and the choice of optimal complexity.

Acknowledgements

A part of this research was made possible through the cooperation with the Royal Danish School of Pharmacy (DFH), Copenhagen, Denmark, which was initiated in 1991 with a study visit of B. Kompare at DFH. Special thanks are due to the Danish Ministry of Education and Research, which granted B. Kompare the 4-month scholarship for studies and research at DFH, during Sept. 1991–Jan. 1992. The same thanks are due also to the Commission of the European Communities, Task Force Human Resources, Education, Training, and Youth for granting B. Kompare an additional grant under TEMPUS Individual Mobility Grants, contract number IMG-91-YU-0053, to successfully accomplish his 4-month retraining period in ecological modelling at the above-mentioned institution. We are obliged to Prof. Giuseppe Bendoricchio and his coworkers at the University of Padua, Department of Industrial Chemistry, for providing us with the data from the Lagoon of Venice, for helping with their expertise, and for being our warm hosts during our stay in Padua, which was realised in the frame of TEMPUS JEP No: 4724-92/1. Two of the coauthors (I. Bratko and S. Džeroski) were partially supported by the Ministry of Science and Technology of The Republic of Slovenia. Aram Karalič and Ljupčo Todorovski (with S. Džeroski) wrote learning programs RETIS and LAGRANGE (respectively) used in the experiments. Doris Ličan carried out some of the experiments with RETIS. Thanks to anonymous reviewers for helpful comments on the manuscript.

References

Beck, M.B., 1978. Is it possible to set up a more general procedure for handling a specific lake modelling system? In: S.E. Jørgensen and D.R.F. Harleman (Editors), Hydrophysical and Ecological Modelling of Deep Lakes and Reservoirs. Summary report of a IIASA workshop, 12–15 December 1977, Laxenburg, Austria. IIASA, Laxenburg, pp. 25–33.

Beck, M.B., 1983. A procedure for modelling. In: G.T. Orlob (Editor), Mathematical Modelling of Water Quality: Streams, Lakes, and Reservoirs. Willey IIASA International Series on Applied Systems Analysis, No. 12, Laxenburg, Austria, pp. 11–41.

Bossel, H., 1992. Real-structure process description as the basis of understanding ecosystems and their development. Ecol. Model., 63: 261–276.

Bratko, I., 1990. PROLOG – Programming for Artificial Intelligence, 2nd ed. Addison Wesley.

Bratko, I., 1993. Applications of machine learning: towards knowledge synthesis. New Generation Comput. J., 11 : 343–360.

Bratko, I., Mozetič, I. and Lavrač, N., 1989. KARDIO: A Study in Deep and Qualitative Knowledge for Expert Systems. MIT Press, Cambridge, MA.

Cestnik, B., Kononenko, I. and Bratko, I., 1987. ASSISTANT 86: A knowledge-elicitation tool for sophisticated users. In: I. Bratko and N. Lavrač (Editors), Progress in Machine Learning. Sigma Press, Wilmslow, pp. 31–45.

Davis, J.R., Nanninga, P.M., Hoare, J.R.L. and Press, A.J., 1989. Transferring scientific knowledge to natural resource managers using artificial intelligence concepts. Ecol. Model., 46: 73–89.

Dejak, C., Franco, D., Pastres, R., Pecenik, G. and Solidoro, C., 1993. An informational approach to model time series of environmental data through negentropy estimation. Ecol. Model., 67: 199–220.

Džeroski, S. and Todorovski, L., 1993. Discovering dynamics. In: P.E. Utgoff (Editor), Machine Learning. Proceedings of the Tenth International Conference. Morgan Kaufmann, San Mateo, CA, pp. 97–103.

Frenzel, L.E., 1987. Crash Course in Artificial Intelligence and Expert Systems. Howard W. Sams & Co., Indianapolis, IN.

Guariso, G., Rizzoli, A. and Werthner, H., 1992. Identification of model structure via qualitative simulation. IEEE Trans. Syst. Man Cybern., 22: 1075–1086.

Jørgensen, S.E., 1986. Fundamentals of Ecological Modelling. Elsevier, Amsterdam.

Jørgensen, S.E., 1992. LAKE Model for IBM PC. International Lake Environment Committee, Otsu, Shiga, Japan.

Karalič, A., 1992. RETIS: A Knowledge Acquisition System. User's Manual. Software version: 2.07/nf. Ljubljana.

Lavrač, N., Urbančič, T. and Filipič, B., 1989. Tehnologija Expertnih Sistemov (The Technology of Expert Systems). University of Ljubljana, Slovenia.

Muetzelfeldt, R., Robertson, D., Bundy, A. and Uschold, M., 1989. The use of prolog for improving the rigour and accessibility of ecological modelling. Ecol. Model., 46: 9–34.

Patten, B.C., 1982. On the quantitative dominance of indirect effects in ecosystems. In: W.K. Lauenroth, G.V. Skogerboe and M. Flug (Editors), Analysis of Ecological Systems: State-of-the-Art in Ecological Modelling. Developments in Environmental Modelling, 5. Elsevier, Amsterdam, pp. 27–37.

Patten, B.C., Higashi, M. and Burns, T.P., 1990. Trophic dynamics in ecosystem networks: significance of cycles and storage. Ecol. Model., 51: 1–28.

Reinhardt, E., Wright, A.H. and Jackson, D.H., 1989. An advisory expert system for designing fire prescriptions. Ecol. Model., 46: 121–133.

Ritchie, J.R., 1989. An expert system for a rangeland simulation model. Ecol. Model., 46: 91–105.

Salski, A., 1993. Fuzzy knowledge-based models in ecological research. Ecol. Model., 63: 103–112.

Starfield, A.M., Farm, B.P. and Taylor, R.H., 1989. A rule-based ecological model for the management of an estuarine lake. Ecol. Model., 46: 107–119.

Šterbáček, Z., Škopek, V. and Zavázal, V., 1990a. A composite landscape ecology prognostic expert system – COLEPES. Part I. System philosophy and design. Ecol. Model., 50: 145–156.

Šterbáček, Z., Pomije, J. and Škopek, V., 1990b. A composite landscape ecology prognostic expert system – COLEPES. Part II. Application of the heuristic model HLEPES to prognoses of forest structures more resistant to immisions and climate change. Ecol. Model., 52: 225–233.

Turban, E. and Watkins, P.R., 1988. Applied Expert Systems. Studies in Management Science and Systems, Vol. 17. North-Holland, Amsterdam.

Young, 1992. Parallel processes in hydrology and water quality: Objective inference from hydrological data. In: R.A. Falconer (Editor), Water Quality Modelling. Ashgate, Hants, UK, pp. 10–52.