Neumann, M. & Baumeister, J. (in press), A rule-based vs. a model-based implementation of the knowledge system LIMPACT and its significance for maintenance and discovery of ecological knowledge, in: Lek, S., Scardi, M., Verdonschot, P. Jorgensen, S.E. & Park, Y.S. (Ed.) Modelling community structure in freshwater ecosystems, Springer-Verlag, Berlin. ??-??

A rule-based vs. a set-covering implementation of the knowledge system LIMPACT and its significance for maintenance and discovery of ecological knowledge

Michael Neumann[1*], Joachim Baumeister[2]

[1]Institute of Ecology, Department of Limnology; Friedrich-Schiller-University of Jena, Dornburgerstr. 159, D-07745 Jena, Germany
[2]Department of Artificial Intelligence and Applied Computer Science; University of Wuerzburg, Am Hubland, D-97074 Wuerzburg, Germany
[*]Author to whom all correspondence should be addressed: Tel: +49-3641-657592; Fax: +49-3641-657592; email: m.neumann@uni-jena.de

Abstract
The knowledge system LIMPACT estimates the pesticide contamination of small lowland streams within agricultural catchment areas. The system considers the abundance of 39 macroinvertebrate taxa during four time frames (T1: March/April, T2: May/June, T3: July/August and T4: September/October) within a year. The four diagnoses Not Detected (ND), Low (L), Moderate (M) and High (H) pesticide contamination represent a calculated annual toxic sum without any specification of the chemical agents. In this paper, we present a new model-based implementation of the existing knowledge system LIMPACT using set-covering relations including diagnosis exclusions. This type of knowledge base outperforms the former rule-based implementation in size and complexity, knowledge acquisition costs and explanatory characteristics. We were able to extract a common and average appearance of taxa in the specific group of streams. A wide range of common taxa with a tendency to more taxa in less severely contaminated streams was observed. Only a few taxa indicate exclusively a specific contamination class. For the exclusion conditions there was a clear trend for more taxa to exclude streams in the High pesticide contamination category than in the other classes.

Keywords
Model-Based, Biological indicator, pesticide contamination, small streams, expert system, evaluation

Introduction
Small streams collectively add up to an enormous length on the landscape level, so that the conservation and protection of their aquatic community should be a major concern. In catchment areas with agricultural activities, these streams are subject to various stressors. During heavy rainfall, runoff from agricultural fields may introduce soil, nutrients, and pesticides and increases discharge (Cooper, 1993; Neumann and Dudgeon, 2002). It has been shown that the impact of pesticides is an important parameter of influence for the aquatic fauna (Liess and Schulz, 1999; Schulz and Liess, 1999). No regular monitoring systems have been established for these agricultural non-point sources of pesticides. Because of its short-term character (Kreuger, 1995), only rainfall event-controlled

sampling methods can reflect such transient pesticide contamination (Liess et al., 1999), which makes its detection via chemical analysis costly.

In this field the use of a biological indicator system brings a number of benefits. The main advantage is its easy, cost-efficient application. When used to monitor toxic contamination, it additionally indicates the ecotoxicological effect of the contaminant. A biological indicator system also provides information on the long-term effects of contamination, whereas information from each chemical measurement applies only to the time the measurement was taken.

There is a wide range of biological indicator systems to evaluate water-quality parameters. In Great Britain RIVPACS (Wright et al., 1998) predicts the macroinvertebrate fauna to be expected at a site in the absence of environmental stress and can be used to evaluate the present fauna. In the Netherlands, a similar approach is used for STOWA (Peeters et al., 1994). In Scotland, the integrated evaluation system SERCON (Boon et al., 1998) and in USA the Rapid Bioassessment Protocols (Resh et al., 1995) were developed. In Germany, the saprobic index is well established to evaluate the biodegradable organic pollution in running waters (Friedrich, 1990). Systems to monitor heavy metals (Wachs, 1998) and acidification (Brakke et al., 1994) have been developed. However, no biological indicator system has yet estimated the pesticide contamination of small streams via benthic macroinvertebrate indicators.

To fill this gap we developed a biological indicator system that estimates the pesticide contamination of small streams. In order to consider the ecological complexity and the uncertain knowledge in this domain, we implemented a diagnostic knowledge system. The advantages are that knowledge systems utilize uncertain expert knowledge and ideally come to the same solution as would the expert. The user has full control over the knowledge system, can scrutinize the solution, and interactively change the question trail. The database and the development of the rule-based knowledge system LIMPACT (from limnology and impact) was presented in (Neumann et al., 2002a; b).

Here, we present a new implementation of the knowledge base using a set-covering approach. The new implementation was motivated by the complex maintainability of the rule base of the former implementation. Thus, it was difficult to extend the knowledge base by rules for new taxa in conjunction with an appropriate scoring of these rules. In contrast to the rule-based approach, set-covering models are intended to minimize the knowledge acquisition costs, since models can be built and extended incrementally in a simple manner. In this paper a performance comparison of the former rule-based and a new set-covering implementation, based on the classification accuracy, the complexity of the knowledge base, and knowledge acquisition costs, is presented. Furthermore, with the new set-covering approach we have been able to extract ecological knowledge about the common appearance of the macroinvertebrate biocoenosis in small, pesticide contaminated streams.

The knowledge system LIMPACT

We developed the knowledge system LIMPACT using the shell-kit D3 (http://www.d3web.de), which is applicable for diagnostic tasks, provides a web-based user interface (d3web) and offers a visual knowledge acquisition environment for a wide range of knowledge types (Puppe, 1998; Puppe et al., 2001). The diagnoses of LIMPACT estimate the pesticide contamination of small streams. They represent a calculated annual toxic sum (for details on types of pesticides see Neumann et al., 2002a) without any specification of the chemical agents. Therefore, the vital diagnoses of LIMPACT are four classes of pesticide contamination named: Not Detected (ND), Low (L), Moderate (M) and High (H) pesticide contamination. The required input parameters (observations) of LIMPACT are abundance data for aquatic macroinvertebrate taxa in a stream. We

established four time frames for which information about abundance is requested. The time frames are T1 (March/April), T2 (May/June), T3 (July/August), and T4 (September/October). LIMPACT allows abundance values to be entered for these four periods of the year for the 39 taxa named in Table 3. Additionally, LIMPACT interprets the increasing or decreasing abundance dynamics of each taxon.
We differentiate between *positive indicator* (PI) taxa, which indicate contamination by high abundance values and positive abundance dynamics, and *negative indicator* (NI) taxa, which exclude contamination and indicate none or low contamination by high abundance values and positive abundance dynamics.
Besides abundance data, LIMPACT evaluates 9 basic water-quality and morphological parameters, such as stream size or conductivity of the water, to characterise a given stream. For simplification, these parameters are abstracted to qualitative values. These abstractions are used for determining the type of stream (for details see Neumann et al., 2002a), because LIMPACT only contains knowledge applicable to small lowland streams within agricultural catchments and cannot make a distinction between pesticides and other types of impact. Hence streams affected by any of the latter factors are excluded to ensure that the impact of pesticide is the main stressor to the aquatic macroinvertebrate fauna. At this stage, such interfering factors include industrial waste impact, severe organic contamination, and extreme chloride or pH values. Additionally, no highland or large streams are considered. The potential application of LIMPACT is for annual monitoring of streams and would reduce costly chemical analysis to the mandatory cases. Furthermore, it could be used to evaluate the success of risk mitigation strategies in the catchment designed to reduce the impact of pesticides. The system is available over the internet via http://www.limpact.de

Methods of knowledge engineering

Ontological knowledge about diagnoses, parameters, and abstractions was used when we implemented the two versions (rule-based and set-covering) of the diagnostic knowledge. In general, diagnostic knowledge states relations between diagnoses and observations and describes how to obtain a diagnosis for a given set of observed parameters. For the acquisition of diagnostic knowledge we have to consider the following aspects and requirements.
As mentioned in Puppe (1998), developing diagnostic knowledge systems is still a time- and cost-intensive task. A variety of knowledge representations have been designed and evaluated to build diagnostic systems effectively, but practical maintenance of such systems is still a difficult issue. In general, we can emphasize the following requirements for a successful knowledge engineering project:

- Understandability of the knowledge representation
  The representation is easily and quickly understood by the domain specialist (expert). This property enables a quick initiation of the development project.

- Incremental development characteristics
  For a rapid development cycle it is helpful to start with extremely simple knowledge, which can be extended incrementally to increase the diagnostic quality.

- Maintainability of the implemented knowledge
  The implemented knowledge base needs to be manageable even if the size of the system increases.

- Explanation facilities
  Furthermore, the representation should allow for the generation of comprehensive explanations to scrutinise the resulting diagnoses.

In the following we present the methods of the two knowledge representations we used for developing two versions of the LIMPACT knowledge system. We also compare their characteristics with respect to their maintainability and reasoning accuracy.

The former rule-based approach

For the first development of the LIMPACT system we applied a heuristic rule-based formalism called *diagnostic scores* (Puppe, 2000) and implemented the rules with the shell-kit D3. Here heuristic classification is based on rules of the following kind:

IF observation OBS_i then give diagnosis D the score Z

The observations OBS_i were clearly defined as the abundances of taxa, whereas the diagnoses are the graded amount of pesticide contamination in the stream, i.e. *Not Detected* (ND), *Low* (L), *Moderate* (M) and *High* (H). The domain expert estimated certain scores (negative or positive) to characterise types of stream contamination on the basis of given abundance data or combinations of them. D3 provides a fixed range of seven positive (P1=+5% to P7 =+100%) and seven negative (N1=-5% to N7 -100%) scores, which has been proven to be useful in various previous applications of D3. Reasoning with scores is easy and understandable for the expert: Given a true condition, the corresponding rule fires and adds the stated score to the specified diagnosis. The sum of two equal categories is aggregated to the next higher category (e.g. P3+P3=P4). A diagnosis about the pesticide contamination is established (confirmed), if the aggregation of the given scores exceeds the category P5.

For a detailed description of the development and evaluation of the rule-based version of LIMPACT we refer to (Neumann et al., 2002a; b). The system has been operational since February 2001 and can be used via the web (http://www.limpact.de).

The new set-covering approach

It has been shown that model-based representations are more appropriate for developing maintainable and explanatory knowledge systems (David et al., 1993). For the development of a model-based approach of LIMPACT we applied set-covering models, which allow for an incremental development of diagnostic systems (Baumeister and Seipel, 2002; Baumeister et al., 2003). Set-covering models describe relations like

Diagnosis D typically covers observation OBS_i.

These relations are called covering relations and we say that OBS_i is covered by diagnosis D. As in the former rule-based implementation of LIMPACT, the diagnoses D were defined by the four different contamination classes, whereas the observations OBS_i are the abundances of taxa.

After implementing simple covering relations for the most typical diagnosis-parameter relations we added *weights* for parameters to the model. With weights we can emphasize that some parameters have a more significant diagnostic importance than other parameters, e.g. parameters stating clear positive indicators. During a second improvement phase we extended the set-covering model by exclusion conditions, which contain knowledge about a categorical exclusion of specific contamination classes (e.g. if we did not find an increasing abundance of a negative indicator taxon in a highly contaminated stream).

Reasoning with set-covering models is very simple: Given a set of observed parameters *OBS*, it uses a simple hypothesize–and–test strategy, which picks a hypothesis *H* (set of diagnoses) in the first step and tests it against the given observations in a second step. The test is defined by calculating a quality measure, which expresses the covering degree of the hypothesis *H* with respect to the observed findings *OBS*. The quality measure q of a hypothesis *H* is defined as follows

$$q(H,OBS)=\frac{\varpi(OBS^{+}_{\mathrm{cov}},H)}{\varpi(OBS^{all}_{\mathrm{cov}},H)+\varpi(OBS^{un\exp l},H)},$$

where $\varpi(OBS^{all}_{\mathrm{cov}},H)$ is the weighted sum of all covered and observed parameters of hypothesis *H* and $\varpi(OBS^{+}_{\mathrm{cov}},H)\subseteq\varpi(OBS^{all}_{\mathrm{cov}},H)$ is the weighted sum of all covered and correctly observed parameters of hypothesis *H*. A parameter is correctly observed if the observed value of the parameter corresponds to the value specified in the covering relation. $\varpi(OBS^{un\exp l},H)$ sums all parameters that are observed but not covered by the hypothesis *H*. Clearly, for a given hypothesis *H* it holds that $OBS = OBS^{all}_{\mathrm{cov}} \cup OBS^{un\exp l}$. A hypothesis is not considered for a given observation if one of its exclusion conditions evaluates true. Besides weights and exclusion conditions, set-covering models can be extended by similarity measures, complex covering relations and constrained covering relations (Baumeister and Seipel, 2002).

Results

Size and complexity

For the implementation of LIMPACT we defined 9 variables (see Neumann et al. 2002a) describing the stream (i.e. structural parameters) and 39 variables representing abundances of different taxa. Each abundance variable can record abundances for the four defined time frames. Furthermore, we specified four diagnoses for the contamination classes of a stream as well as a diagnosis for detecting unsuitable streams. This ontological knowledge was augmented by diagnostic knowledge represented either by heuristic rules (former approach) or by set-covering relations (new approach).
The former rule-based version of LIMPACT contains 921 diagnostic rules (see Table 1) with scores to establish or to de–establish a diagnosis. Diagnostic rules are of the following kind:

IF (Rule Condition C) THEN give diagnosis D score S.

The complexity of these rules is moderate, which means that the rule condition mostly contains between two and four combined single conditions connected by Boolean operators (e.g. and, or, not). A single condition evaluates whether a taxon’s abundance is above a given threshold, i.e. a single observation. Additionally, for each rule an appropriate diagnosis score was defined by the expert.
For the set-covering knowledge base, we implemented 816 simple covering relations (see Table 1) of the following kind:

Diagnosis D covers the observation of taxon T with abundance A.

We can see that these relations are simpler than the implemented rules described above. In contrast to the rule-base, we only consider one taxon’s abundance information, disregarding other taxa also covered by the same diagnosis. We also do not consider scores for diagnoses.
Table 1 gives the complexities of the implemented knowledge in more detail. Whereas the last column shows the number of set-covering relations for each contamination class, we extended the presentation for the rule-based version. Thus, we depict the overall number of diagnostic rules besides the number of covering relations, and display more precisely the number of rules with 1 to 7 single conditions in the first columns of the table. Rules with 2-4 conditions dominate the rule-based version.

Table 1: Size of the two implemented rule-based and set-covering knowledge bases for each diagnosis. The left side of the table shows the complexity of the rule conditions for the rule-based approach in more detail. The last two columns give an overview of the size of the rule-based and the set-covering knowledge base.

Contamination	Rule-Based Knowledge Base								Set-Covering Knowledge Base
	Number of evaluatable symptoms in rule condition							Total	Total
	1	2	3	4	5	6	7		
Not Detected	0	113	82	39	13	3	1	251	212
Low	0	85	75	38	7	4	1	210	202
Moderate	1	105	75	44	5	2	2	234	206
High	1	112	76	28	8	1	0	226	195
Sum	3	417	311	153	38	16	11	921	815

Knowledge acquisition costs

Comparing the size of the two knowledge bases, Table 1 shows that the number of implemented knowledge elements is comparable. The size of the set-covering knowledge base is even a little bit smaller. These characteristics are illustrated by the fact that the expert required about six weeks to implement the rule-based version of LIMPACT versus two weeks for implementing the set-covering counterpart.

Table 2: Result of the classification of 146 investigations per stream and year using the rule-based (RB) and set-covering (SC) implementation of LIMPACT. The measured real contamination is given according to the four classes and compared with the percentage of cases classified by LIMPACT into the four classes plus not classified. Correct classifications are indicated by bold values. The number of cases per contamination class is given in parentheses.

real contamination	classification result (%)									
	Not Detected		Low		Moderate		High		not classified	
	RB	SC	RB	SC	RB	SC	RB	SC	RB	SC
Not Detected	**90.4**	**96.2**	0	0	1.9	0	0	0	7.7	3.8
(52)	(47)	(50)	(-)	(-)	(1)	(-)	(-)	(-)	(4)	(2)
Low	16.7	0	**80.0**	**93.3**	0	6.7	0	0	3.3	0
(30)	(5)	(-)	(24)	(28)	(-)	(2)	(-)	(-)	(1)	(-)
Moderate	2.5	0	0	2.5	**72.5**	**87.5**	7.5	0	17.5	10
(40)	(1)	(-)	(-)	(1)	(29)	(35)	(3)	(-)	(7)	(-)
High	0	0	0	4.2	0	12.5	**87.5**	**79.1**	12.5	4.2
(24)	(-)	(-)	(-)	(1)	(-)	(3)	(21)	(19)	(3)	(1)

Classification results

The classification result of both rule-based and set-covering implementation was calculated with the same cases that were used to develop the system. This was necessary because no independently obtained stream investigations, including macroinvertebrate abundance data and chemical pesticide measurements, were available.

A detailed evaluation is presented by Neumann et al. (2002b) for the rule-based (RB in Table 2) implementation. For RB Table 2 shows that the correct diagnosis of the 146 cases is established by LIMPACT in 72.5 to 90.4% of the cases, with better results for uncontaminated sites. The evaluation showed a very good classification result. Most errors occur between ND and Low and on the other hand between Moderate and High contamination. A high percentage of cases were not classified. Because of our conservative

approach, LIMPACT established no diagnosis instead of a wrong one for cases with insufficient data availability. Possible reasons for classification errors and not classified cases can be related to uncertainty in the sampling and identification methods and the number of sampling dates within a year. The more data the user provides, the more rules can be activated. Consequently, the chance of a correct classification increases.
For the set-covering (SC in Table 2) implementation Table 2 gives the classification result for the same 146 cases as for the rule-based approach. The correct diagnosis is found in 79.1 to 96.2 cases, which is a better classification result than for the rule-based implementation. Only the highly contaminated cases show a decline in classification result and at the same time an increase in wrong classifications. Additional errors occur between the Low and Moderate contamination classes.

Explanatory characteristics

The two implementations differ not only in the way the knowledge is represented but also in the way new knowledge can be extracted and discovered from the knowledge bases. For the rule-based implementation the domain expert found it difficult to gain any new insights. The explanatory characteristics are complex because the knowledge is represented in small pieces (rules) and is weighted with different scores. In the following, we give only four rules as example:

- IF Agabus at T2 in [2; 9] THEN Contamination High P3
- IF Agabus at T2 > 9 THEN Contamination High N4
- IF Anabolia at T1 in [0; 80] THEN Contamination High P2
- IF Anabolia at T1 > 80 THEN Contamination High N3

The different scores to establish (here: P2 and P3) or to de-establish (here: N3 and N4) a diagnosis (here: High) make it difficult to obtain a general overview. To extract from the rule-based knowledge base how the aquatic community of an average stream with e.g. High pesticide contamination appears, the domain expert has to interpret the rule for and against the High diagnosis and has to interpret the different scores.
The set-covering implementation has a better explanatory characteristic, because of its more straightforward design. In the following, we give only two covering relations as example:

- Contamination High: Agabus at T2 in [2; 9]
  Anabolia at T1 in [0; 80]

Each covering relation represents a characteristic of the considered contamination class (here High). The domain expert simply looks at all covering relations of one specific diagnosis and gains an overview. The same is true for exclusion conditions. They represent those characteristics that are not the case for the considered contamination class.

Table 3: The 39 indicator taxa of the knowledge system LIMPACT and the type of indicator (N = negative; P = positive indicator). For the covering relations, taxa with more than 30% appearance in the considered class are marked and for the exclusions, taxa which activated more than 10% of the sum of the rest of the classes are marked.

			Covering Relations				Exclusions			
Order	Taxon	Type of Indicator	ND	L	M	H	not ND	not L	not M	not H
Turbellaria	*Dugesia gonocephala*	N	X	X						X
Oligochaeta	*Erpobdella octoculata*	P	X	X	X	X				
	Glossiphonia complanata	N	X	X	X	X				

	Glossiphonia heteroclita	P	X		X	X				
	Tubificidae	P	X	X	X	X	X			
	other Oligochaeta	N	X							X
Gastropoda	*Pisidium* sp.	N	X	X	X	X				X
	Potamopyrgus antipodarum	P				X	X			
	Radix ovata	P	X	X	X	X				
Amphipoda	*Gammarus pulex*	P	X	X	X	X				
Isopoda	*Asellus aquaticus*	N	X	X	X				X	
Plecoptera	*Nemoura cinerea*	N								X
Coleoptera	Dytiscidae	N	X							X
	Agabus sp.	N							X	
	Platambus maculatus	N								
	Elmis sp.	N		X	X					
	Haliplus sp.	N								
	Helodes sp.	N	X	X	X				X	X
Diptera	Ceratopogonidae	P			X	X				
	Chironomidae "white"	N	X		X	X				X
	Chironomidae "red"	N	X	X	X	X			X	X
	Limoniidae	N		X	X					X
	Ptychopteridae	N								X
	Simuliidae	N	X	X	X					X
	Tipulidae	N						X		
	Other Diptera	N	X							X
Ephemeroptera:	*Baetis vernus*	N			X					X
	Baetis sp.	N			X					
	Ephemera danica	N								X
Megaloptera:	*Sialis lutaria*	N		X	X					X
Trichoptera:	*Hydropsyche angustipennis*	N								
	Anabolia nervosa	N			X					
	Chaetopteryx villosa	N	X	X						X
	Halesus radiatus/digitatus	P								
	Ironoquia dubia	P	X							
	Limnephilus lunatus	N						X		
	Limnephilus extricatus	N						X		
	Limnephilus rhombicus	N								
	Plectrocnemia conspersa	N								

Discovery of ecological knowledge

Using the set-covering implementation we were able to discover the common macroinvertebrate community of an average stream. For each of the four diagnosis classes, we analysed which covering relation and which exclusion was activated most frequently. For the covering relation we considered those activated in more than 30% of the cases within the contamination class and for the exclusion we considered those activated in more than 10% of the sum of the rest of the classes. Table 3 indicates which taxa activated the most covering relations and exclusions. For the sake of simplicity we do not indicate abundance values and do not itemise each single covering relation. Generally speaking, the type of the indicator specifies whether the taxon is found in higher abundance in more highly contaminated streams (positive indicator) or in uncontaminated streams (negative indicator). Bearing in mind all this information, Table 3 illustrates a theoretical average community in the four contamination classes.

As Table 1 shows, we implemented only 8% (212 to 195) fewer covering relations for the High contamination class vs. the ND class. Table 3 shows that these covering relations are activated by 35% (17 to 11) fewer common taxa in the High contamination class than in the ND class. This indicates that in highly contaminated streams fewer taxa are common. At the same time the large number of exclusions indicates that in this contamination class

16 taxa cannot be found with high abundances. Most common taxa are found in streams classified as Moderate, which may indicate highly variable conditions in this type of stream.

The analysis shows that considering the appearance of the common taxa, the stream classes look very similar. Eleven taxa appear at least in three diagnoses classes, separated by the abundance only. Four taxa clearly indicate the ND class (e.g. Oligochaeta, Dytiscidae), but none the H class. Only a few taxa appear in the ND and/or L classes and exclude the H class (e.g. *Dugesia gonocephala*, Oligochaeta) and only *Potamopyrgus antipodarum* indicates the H class and excludes the ND class. Some taxa indicate a specific class by their low abundance and exclude the same class by high abundances (e.g. Tubificidae, *Pisidium sp.*). Overall, we found a wide range of common taxa with a tendency towards more taxa in less severely contaminated streams. For the exclusion conditions a clear trend, with more taxa excluding the more highly contaminated streams, was likewise found.

Discussion

Size and complexity

The reduction of size and complexity of knowledge bases is the main focus of knowledge engineering research. Both aspects are crucial for developing and maintaining successful knowledge systems. It has been shown that knowledge bases tend to be confusing and unmanageable if their size increases and the complexity of the embedded knowledge develops excessively.

Comparison of the knowledge bases presented here shows that the number of covering relations in the set-covering approach is only slightly smaller than the number of implemented rules in our rule-based system. Nevertheless, the complexity of the modeled set-covering knowledge is significantly simpler than the implemented rules-based knowledge. When adding rules for taxa to the rule base, we also have to consider the associated diagnosis scores. These scores interact with other rules deriving the same diagnosis and therefore have to be obtained by thorough analysis. Thus, adding a new rule to the knowledge base can demand reconsideration of all rules (and of the associated scores) deriving the same diagnosis. In contrast to these interwoven rules, set-covering relations can be viewed as isolated knowledge elements without mutual interdependencies. For a new taxon we only have to define relevant covering relations for the four diagnoses, i.e. contamination classes, and the new taxon. In general, this means that we have to define the abundance of the new taxon for each diagnosis, if we expect the taxon to occur with the given diagnosis. If available, we can additionally define abundance trends (positive or negative) between the time frames T1, T2, T3,T4 for the new taxon and each diagnosis.

Knowledge acquisition costs

The costs of knowledge acquisition often can be measured only by the time the domain specialist (expert) or engineer had spent in developing the knowledge system. For maintenance purposes we also need to consider the time the developer needs to change or extend the knowledge base. In our experiences with LIMPACT, the modular characteristics of the set-covering relations had a direct impact on knowledge acquisition costs. The expert found the set-covering representation easy to understand and to apply to the diagnosis problem. In contrast to the rule-based version of LIMPACT, he did not need to consider the interconnections between rules deriving the same diagnosis. This experience is emphasized by the time the expert spent to develop the two knowledge bases: implementing the rule-based knowledge took about 6 weeks vs. 2 weeks for defining the set-covering model.

Classification results

It is obvious that the classification accuracy of a diagnostic system is the key factor for its user acceptance. A user is more likely to accept that the system cannot supply a diagnosis for a particular case, but will lose confidence if the system derives wrong diagnoses in some cases. For this reason, a system should not only provide a solution for a given case, but furthermore should deliver a “confidence level” for the diagnosis that is obtained. This confidence level can depend on the score of the diagnosis or an overall “believability” function defined by the developer of the knowledge base.

As described in the previous section, the classification accuracies of the rule-based and the set-covering system are comparable. Nevertheless, the diagnostic system applying set-covering knowledge outperforms the rule-based version for contamination classes Not Detected, Low, and Moderate. The rule-based implementation only outperforms the set-covering implementation for highly contaminated streams. One can say the rule-based version of LIMPACT has no (high) confidence level for streams with contamination classes Low and Moderate, while the set-covering implementation has a lower confidence level for the diagnosis of highly contaminated streams. Reasons include the fact that we have not implemented any covering relations interpreting the absence or the decreasing abundance dynamic, which could indicate highly contaminated streams. For the domain expert the absence of a taxon or its decreasing abundance is difficult to interpret, because the causal connection explicitly to pesticide contamination is uncertain.

Explanatory characteristics

The set-covering knowledge base is much more suitable for discovering ecological knowledge than the rule-based implementation. The covering relation and the exclusions can be easily interpreted as characteristics of the group of streams considered. By analysing the frequently used relations we found the common taxa for each contamination class. This procedure was simple and fast. For the rule-based knowledge base this would have been a time-consuming process, because of the interpretation of the rules and the scores.

Other knowledge representations, such as case-based reasoning, also cause problems in finding common and average characteristics of the considered diagnoses classes. For implementation they represent a set of characteristics at the same time and therefore cannot activate each characteristic separately. In summary, we can say that the model-based knowledge representation using a set-covering interpretation is easy to implement. It outperforms the rule-based implementation in size, complexity, and maintainability and helped the domain expert to discover new ecological knowledge at a higher level.

Acknowledgements

The former rule-based implementation of LIMPACT was financed by the Deutsche Bundesstiftung Umwelt (German Federal Environmental Foundation), An der Bornau 2 in 49090 Osnabrück, Germany and was supported by Matthias Liess and Ralf Schulz of the Zoological Institute at the Technical University of Braunschweig, Germany. The authors would like to thank Frank Puppe and Dietmar Seipel of the Department of Computer Science at the University of Würzburg, Germany for their helpful discussions and comments.

References

Baumeister, J. and Seipel, D., 2002. Diagnostic reasoning with multilevel set-covering models. Proceedings 13th. International Workshop on Principles of Diagnosis (DX-02), Semmering (Austria). 58-64.

Baumeister, J.; Seipel, D.; Puppe, F., 2003. Incremental development of diagnostic set-covering models with therapy effects. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems (IJFUKS), in press
Boon, P.J., Holmes, N.T.H., Maitland, P.S., Rowell, T.A. and Davies, J., (1998). A system for evaluating rivers for conservation (SERCON): development, structure and function. In Boon, P.J. and Howell, D.L. (Editors), Freshwater Quality: Defining the undefinable?, The Stationary Office, Edingburgh, 299-326.
Brakke, D.F., Baker, J.P., Böhmer, J., Hartmann, A., Havas, M., Jenkins, A., Kelly, C., Ormerod, S.J., Paces, T., Putz, R., Rosseland, B.O., Schindler, D.W., and Segner, H., 1994. Group report: Physiological and ecological effects of acidification on aquatic biota. In: C.E.W. Steinberg and R.F. Wright (Editors), Acidification of freshwater ecosystems: Implications for the future, John Wiley and Sons Ltd., Vol. 1: 275-312.
Cooper, C.M., 1993. Biological effects of agriculturally derived surface-water pollutants on aquatic systems - a review. Journal of Environmental Quality, 22: 402-408.
David, J.-M., Krivine, J.-P., and Simmons, R., 1993. Second-Generation Expert Systems Springer.
Friedrich, G., 1990. Eine Revision des Saprobiensystems. Zeitschrift für Wasser- und Abwasserforschung, 23: 141-152.
Kreuger, J., 1995. Monitoring of pesticides in subsurface and surface water within an agricultural catchment in southern Sweden. British Crop Protection Council Monograph No. 62: Pesticide Movement to Water, 81-86.
Liess, M. and Schulz, R., 1999. Linking insecticide contamination and population response in an agricultural stream. Environmental Toxicology and Chemistry, 18: 1948-1955.
Liess, M., Schulz, R., Liess, M.H.-D., Rother, B., and Kreuzig, R., 1999. Determination of insecticide contamination in agricultural headwater streams. Water Research, 33: 239-247.
Neumann, M., Liess, M., and Schulz, R., 2002a An expert system to estimate the pesticide contamination of small streams using benthic macroinvertebrate as bioindicators, Part 1: The database of LIMPACT. Ecological Indicators, 2 (4): 379-389.
Neumann, M., Baumeister, J., Liess, M., and Schulz, R., 2002b. An expert system to estimate the pesticide contamination of small streams using benthic macroinvertebrate as bioindicators, Part 2: The knowledge base of LIMPACT. Ecological Indicators, 2 (4): 391-401.
Neumann, M. and Dudgeon, D., 2002: The impact of agricultural runoff on stream benthos in Hong Kong, China. Water Research, 36, 3103-3109.
Peeters, E.T.M.H., Gardeniers, J.J.P., and Tolkamp, H.H., 1994: New methods to assess the ecological status of surface waters in the Netherlands. Verhandlungen der Internationalen Vereinigung für Limnologie, 25: 1914-1916.
Puppe, F., 1998. Knowledge reuse among diagnostic problem solving methods in the shell-kit D3. International Journal of Human-Computer Studies, 49: 627-649.
Puppe, F., 2000. Knowledge Formalization Patterns. Proceedings of PKAW in Sydney, Australia.
Puppe, F., Ziegler, S., Martin, U., and Hupp, J., 2001. Wissensbasierte Diagnosesysteme im Service-Support. Konzepte und Erfahrungen Springer Verlag.
Resh, V.H., Norris, R.N., and Barbour, M.T., 1995. Design and implementation of rapid assessment approaches for water resource monitoring using benthic macroinvertebrates. Australian Journal of Ecology, 20: 108-121.
Schulz, R. and Liess, M., 1999. A field study of the effects of agriculturally derived insecticide input on stream macroinvertebrate dynamics. Aquatic Toxicology, 46: 155-176.
Wachs, B., 1998. A qualitative classification for the evaluation of the heavy metal contamination in river ecosystems. Verhandlungen der Internationalen Vereinigung für Limnologie, 26: 1289-1294.

Wright, J.F., Furse, M.T., and Moss, D., 1998. River classification using invertebrates: RIVPACS applications. Aquatic Conservation Marine and Freshwater, 8: 617-631.