# NAIA: A Decision Support System for Predicting E
# from Existing Land Resource Data

**Jan A. Mulder**
Department of Advanced Computing and Engineering
Alberta Research Council
Calgary AB.
Canada T2E 7H7

**Ian G.W. Corns**
Northern Forestry Centre
Forestry Canada
Edmonton AB.
Canada T6H 3S5



## Abstract

The NAIA program is concerned with the design and implementation of an ecologically-oriented spatial and knowledge-based framework to support forest and land resource management. As part of this program a decision support system was designed and implemented with the capability of representing the knowledge used by a forest ecologist to infer a forest ecosystem from a variety of data sources. The system is object-oriented. It has been designed as a classification shell with the capability of representing uncertainty in a hierarchically structured knowledge base. The classification process implements a combination of symbolic and evidential reasoning and it predicts ecosystems from topography, forest cover, and soil maps. The shell operates in conjunction with a GIS. Polygons are at first interpreted as unspecialized eco-units. As features from different maps become available this interpretation is gradually specialized into one or more ecoregions or ecosystem associations. The final result of the inference process consists of a listing of possible ecosystem associations each with an associated belief value. This approach is referred to as *Evidential Discrimination Reasoning*. Three different ecosystem classification systems have been implemented covering several regions of Alberta and Manitoba. Test results indicate a prediction accuracy of 63 - 94%. Two different update rules were experimented with: Weighted Averaging and Dempster/Shafer. Test results favor Weighted Averaging. The decision support system can be used as a practical tool that can assist with the creation of an ecosystem-based land inventory while saving costs and at the same time improving the quality of the classification process.

## 1. Introduction.

This paper has three objectives:

1. To explain the concept of Evidential Discrimination Reasoning and its implementation as a classification shell.
2. To show how the shell can be used in the design of a decision support system capable of inferring ecosystems from topography, forest, and soil data.
3. To show how such a system can be used as a cost effective tool in the creation of an ecosystem-

based land inventory while at the same time improving the quality of the classification process.

The NAIA program is a joint venture between Hughes Aircraft of Canada's Spatial Data Systems division and the Alberta Research Council. Its objective is to provide governments, forest industries, and other resource information users with practical, commercially-supported, and ecologically-oriented decision support tools for forestland planning. Designing the software and working closely with the forestry industry is a multi-disciplinary team of specialists in natural resources and information system technology. This project team is building a series of software tools using Geographic Information Systems, Knowledge-Based Systems, and computer modelling.

The first product produced by the joint venture and also the focal point of this paper is a *classification* shell with the capability to solve a wide variety of classification problems. With NAIA's focus on forestland planning the shell has been used to implement a decision support system with the capability to predict ecosystems from existing map data.

Most of the existing forest inventories in Canada consist of topographic data (e.g. elevation, slope, aspect), forest cover data (mostly commercial tree species), and soil data. Due to political and socio-economic pressure forest management is quickly evolving into forest*land* management guided by the principle of sustainable development and integrated resource management (Jones, 1993). This change is causing an increased pressure on forest companies and governments to turn forestland inventories into ecologically-oriented land inventories. Such inventories must not only contain forest overstory data, but also data about the understory, shrubs, herbs, and other soil and land features. In addition, volume-related information and data about the spatial distribution of the vegetation species described is needed. The latter information becomes particularly important when wildlife habitat is a consideration (Buech et al., 1990).

Some standards for ecological land classification have been introduced in Canada both federally (Ironside, 1980) and provincially (Mitchell, 1981). Most land classification systems take a stratified approach to the classification process. Ecosystems are represented at different levels of specialization (i.e. refinement). An increasing specialization reflects a transition from macro-bioclimate to micro-bioclimate. The Canadian ecological land classification system recognizes several levels of specialization. At increasing levels of specialization these are: *ecozone, ecoregion, ecosite,* and *ecoelement*. In this paper, we will follow the nomenclature used in Corns and Annas (1986) which refers to ecosites as *ecosystem associations*. The level of refinement provided by ecozone and ecoregion level classifications may suffice for some forms of strategic forestland management. At the operational level, however, forestland management requires ecosystem classifications at the ecoregion and ecosystem association level. Classification systems providing this level of specialization exist only for some parts of the country. The NAIA project has used three such classification systems for designated areas in Alberta and Manitoba (Corns and Annas, 1986; Knapik et al., 1988; Knapik et al., 1989).

As mentioned before, many forest companies and governments are currently faced with the problem of how to transform existing land inventories into ecologically-oriented inventories that can support both strategic and operational management. This requires the availability of ecosystem classifications both at a small scale (i.e. ecoregion) and large scale, site specific (i.e. ecosystem association) level. The most straightforward approach to the creation of a new ecosystem inventory is to hire a team of forest ecologists who will use an ecosystem classification system to map the area at a scale of 1 : 15,000 by means of systematic ground observations and air photo interpretation. However, extensive field work and data digitization is a costly practice, especially when forest management areas are very large. A more cost effective approach which is described in this paper is to use an ecosystem prediction decision support system to facilitate the ground truthing process.

The NAIA classification system is a knowledge-based, object-oriented system implemented as a shell. Within this shell ecosystem classes are represented as *models* organized in a specialization hierarchy. The unit of interpretation is a *polygon*, which is obtained from a Geographic Information System. Each polygon is interpreted in terms of one or more ecosystem classes each of which is embedded in a specialization hierarchy. The set of possible ecosystem classes is constrained by *features* which are associated with each polygon. Feature values constrain ecosystem classes in both a symbolic and probabilistic manner. These constraints are represented in the form of *mass functions*. For each polygon the classification process creates a listing of possible ecosystem classes each with an associated belief value. The classification process can be characterized as a process that at first classifies each polygon as an unspecialized *eco-unit*. As different features are introduced this classification is gradually specialized into one or more ecosystem classes located further down in the specialization hierarchy while updating the beliefs of these classes. This approach is referred to as *Evidential Discrimination Reasoning*.

The remainder of this paper is laid out as follows. In section 2 we explain the design of a classification shell that embodies the Evidential Discrimination Reasoning (EDR) concept. Section 3 describes the implementation of EDR as a decision support system (DSS) for predicting ecosystems. Section 4 looks at the results of validating the ecosystem prediction DSS for three different areas in Alberta and Manitoba. As well, this section discusses how the DSS can be used in practice and how it can save costs. The paper is summarized in section 5, with acknowledgments and references in sections 6 and 7 respectively.

## 2. The classification shell

The classification shell is a generic tool for performing classification tasks. The shell consists of two components: a knowledge base component and a process component.

### 2.1 The Knowledge Base.

The knowledge base is object-oriented and it recognizes three types of objects: *primitives*, *models*, and *features*. Primitives and features are derived from a data domain, whereas the models constitute an interpretation domain. Primitives are the unit of interpretation. They are interpreted by models. Both primitives and models can be characterized by means of a property list which takes the form of an attribute-value list.

Primitives and models are linked by means of features. A feature is an entity that is derivable from input data. A feature can be discrete or continuous and can take on a range of possible values. A feature links a primitive with a model by means of a *mass function*. In the classification shell a mass function is a continuous or discrete function that can assume any value between -1 and +1. Mass functions are obtained through knowledge elicitation sessions with domain experts. A mass function expresses the subjective belief of a domain expert in the occurrence of a model for each possible value of the feature represented by the function. Hereby -1 expresses strong disbelief, 0 expresses neutrality, whereas +1 expresses strong belief. The function can be completely or partially defined over the feature's value range. Through the mass function the value of each feature constrains both the models possible and the belief values these models can assume.

The notion of a mass function is based on the idea of obtaining degrees of belief for one question from subjective probabilities for a related question. We are avoiding the term "belief functions" here, because these are usually associated with the Dempster/Shafer theory of evidence (Shafer, 1976). In this paper we want to make an explicit distinction between the *representation* of beliefs, and the *update rules* used to update beliefs. The theory of belief functions combines the definition of belief functions with Dempster's rule for combining degrees of belief (Shafer, 1990).

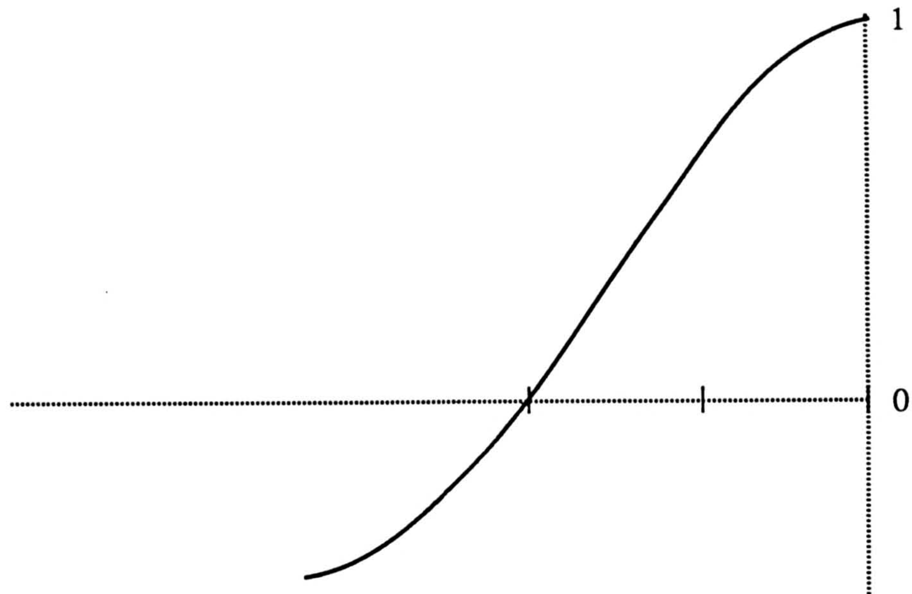


| | | | | | | |
|---|---|---|---|---|---|---|
| **Attribute value:** | **V1** | **V2** | **V3** | **V4** | **V5** | **V6** |
| **Function value:** | ***** | ***** | **- 0.5** | **0** | **0.6** | **1.0** |

Figure 1: an example of a mass function

Our mass functions need not be completely defined over the range of the feature's possible values and this is done for a very particular reason. In constraint-based reasoning problems we often need to combine symbolic and numerical constraints. Partially defined mass functions enable such combinations. Fig. 1 illustrates a mass function that represents the relationship between a feature and a model. For the range of values for which the function is defined, the model is not only permitted to exist, but is also supported by a belief value. The absence of belief values for a subrange of feature values expresses a symbolic constraint. It means that in the range given the model cannot exist. Thus, mass functions have the capability of expressing both symbolic and numerical constraints.

Models can often be organized in the form of an and/or graph more commonly referred to as a *specialization hierarchy*. Such hierarchies represent classes of models with similar properties. The leaf nodes of the hierarchy represent elementary models which describe a primitive unambiguously. All other nodes represent abstract classes of such models. Each class of models inherits the attributes and values of its parent in the hierarchy. When used as an interpretation in the classification process an abstract class represents the fact that all of its descendents are a possible interpretation. The prime representational advantage of specialization hierarchies is their inherent capability to represent primitives at different levels of abstraction or refinement. The prime processing advantage of these hierarchies is their potential for increasing the processing efficiency of the classification process. We will discuss this in more detail in section 2.2.

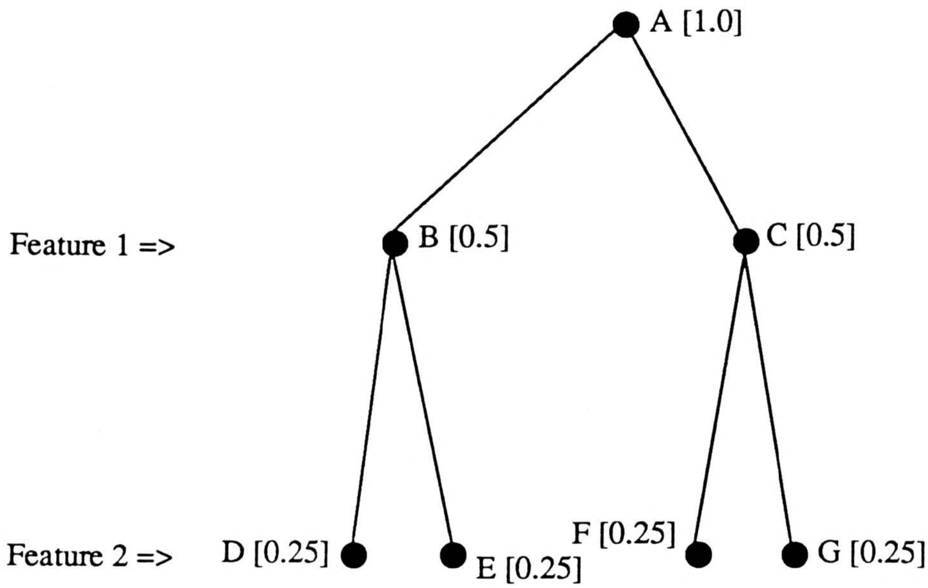


Figure 2: an example of a specialization hierarchy

Fig. 2 illustrates a two-level specialization hierarchy. The role of inheritance in such hierarchies has received a great deal of attention in the Artificial Intelligence Knowledge Representation literature (e.g. Hanson et al., 1978; Brooks, 1981; Brachman, 1982; Mulder, 1988). However, the focus has mostly been on the inheritance of symbolic properties. Although some attention has been given to the propagation of beliefs in hierarchies (Pearl, 1986; Shafer and Logan, 1987) very little attention has been paid to the inheritance of such numerical properties. The root node (A) in Fig. 2 intrinsically represents the fact that each of its successors constitutes a possible interpretation. At the belief level, however, things are less evident. Since A is the root node and is representing the whole interpretation domain we will have to believe A with absolute certainty (i.e. 1). However, we have no information about the distribution of A's belief among its successors. Intuitively, there are at least two possible ways out of this dilemma and we will refer to these as the *strong inheritance* and *weak inheritance* approach.

In the strong inheritance approach each model automatically inherits the mass functions assigned to its parent in the hierarchy. If a specialization hierarchy takes the form of a directed acyclic graph and a model has more than one parent, then rules for parent dominance must be provided. In Fig. 2 B and C inherit A's belief, D and E inherit B's belief and so forth. If, in addition, we adopt a normative approach (i.e. the belief in a model equals the sum of beliefs of its successors), then B and C will obtain 0.5 as a belief, whereas D,E,F, and G obtain 0.25 each. A weak inheritance approach, on the other hand, prevents us from making any inferences about the beliefs of a successor, although, if we adopt the normative approach, we can still ascertain that, for example, the belief in B must be equal to the sum of the beliefs of its successors. However, knowing the belief in B, does not allow for any inferences of the belief in D and E. The classification shell implements the strong inheritance approach.

### 2.2 The Classification Process.

The objective of the classification process is the creation of a listing of possible models for each primitive whereby each model carries an associated belief value. The classification process is based on two principles: the principle of *least commitment* and the principle of *graceful degradation*. The first principle requires that at any time input data (represented as features) are classified at a level of specialization that is appropriate for those data. The second principle implies that the classification process reflects the reliability and availability of input data. More in particular, the classification process must still produce a listing of models if one or more features have no value, even though such results may be less reliable. For this purpose the shell maintains a *confidence value* (a number between 0 and 1) for every model listing.

The least commitment principle is also illustrated in Fig. 2. Each primitive is constrained by one or more features. The belief functions of these features normally constrain classes of models at one particular level of specialization. Feature 1, for example, constrains B and C, whereas feature 2 constrains the models D, E, F, and G. At the start of the classification process each primitive is depicted by the model at the root node of the specialization hierarchy with a belief 1. No features have been considered at this time and the confidence in the current listing will therefore be 0. For each primitive we sequentially process each of its features. Every feature introduced will replace the current model listing by one defined at the level of specialization imposed by the feature. Each feature also has an associated confidence. As more and more features are processed, the model listing becomes more specialized while their belief values are updated. As well, the confidence in the listing increases with the number of features processed. The evidence-based process of specializing a root model into one or more of its specializations is a process of discrimination and is therefore referred to as *Evidential Discrimination Reasoning*.

The selection of a mass function and update rule is a task that must be dealt with in the context of the application (Shafer, 1990). Although the value of a mass function is expected to range between -1 and +1, there are no further restrictions on the use of mass functions and update rules. Mass functions and update rules are therefore represented independently. An explicit distinction is thus made between the *representation* of beliefs and the process that *updates* beliefs.

The processing of each feature takes place in three stages:
1. Hierarchical (symbolic) constraint propagation.
2. Hierarchical belief distribution.

3. Belief updating.

Hierarchical constraint propagation applies the constraints imposed by a new feature to the existing listing of models. It removes the models from the current listing that are inconsistent with the models proposed by the new feature. Conversely, hierarchical constraint propagation also removes those models proposed by the new feature that are inconsistent with the models in the current listing. Hierarchical constraint propagation has been implemented as a network consistency algorithm (Mackworth, 1977). A special version of such an algorithm was developed for hierarchically structured domains (Mulder,1988). A particular advantage of the hierarchical algorithm is its efficient behavior both in space and time. The reason for this behavior is that the size of the model listing (i.e. the domain size) can be suppressed by replacing sets of elementary models by a smaller set of abstract models. Whereas most network consistency algorithms are exponential with respect to the domain size, the hierarchical algorithm is cubic in the domain size in the worst case, but only logarithmic in the domain size in the best case (Mackworth et al., 1985).

Hierarchical belief distribution projects beliefs from one level of specialization onto another. The general assumption is made, that a belief update rule requires two model listings of identical models with similar or different beliefs. Before the beliefs of the models in the current listing can be updated with the beliefs associated with the model listing of a new feature, the beliefs in the current listing and the new listing must be "levelled". This process is best explained by means of an example. Suppose the current model listing is: [B 1.0] (see Fig. 2). Let us also assume that the listing provided by the new feature is: [D 0.7 E 0.3]. Before belief updating can take place, the belief of B must be brought down (i.e. levelled) to the level of D and E. With strong inheritance and a normative approach we can replace [B 1.0] by [D 0.5 E 0.5]. This creates two listings with identical models but with different beliefs.

The choice of an appropriate belief updating rule is dependent on a variety of factors. These include: the domain of application, the appropriateness of a normative approach, and the interdependence of features. As mentioned before, the classification shell itself does not impose any particular update rule. The user can make that choice according to the requirements of the domain of implementation. A mathematical survey of updating rules has been provided by Planchet (1989). A review of the practical meaning of these rules can be found in Smets (1993).

## 3. Implementation.

Ecosystems can be organized as a taxonomy based on commonalities in bio-geoclimatic properties. Different applications require different levels of refinement. At the upper levels in a land classification hierarchy ecoregions provide a framework for small scale planning and land use policy. Examples include the reservation of national parks or protected areas, areas climatically suitable for wilderness creation, forestry, or arable agriculture. Ecoregions are typically mapped at scales of 1 : 2,000,000 to 1 : 250,000. At the lower levels of the hierarchy, ecosites or ecosystem associations provide a framework for large scale land use planning such as roads and access, harvest scheduling, location of environmentally sensitive sites, gravel deposits, wildlife habitat etc. Ecosites or ecosystem associations are typically mapped at scales of 1 : 50,000 to 1 : 12,000.

Several ecosystem knowledge bases have been implemented in the classification shell. The construction of these knowledge bases was the result of an intensive knowledge elicitation process with forest ecologists. Local forest ecosystem knowledge bases have been implemented by using the classification systems for West-Central Alberta (Corns and Annas, 1986), and for the Duck Mountain (Knapik et al., 1988) and Sandilands area (Knapik et al., 1989) in Manitoba. As all knowledge bases are similar in set up, we will only discuss the Duck Mountain area here. The specialization hierarchy for ecosystems is two levels deep. The root node is an eco-unit (level 0), which can be specialized into a series of ecoregions (level 1). Ecoregions, in turn, can be specialized into ecosystem associations (level 2). Fig.3 illustrates the Duck Mountain specialization hierarchy. Duck Mountain has two ecoregions: Boreal Mixed Woods (BMW) and Mid Boreal Lowlands (MBL). BMW has 13 specializations: BMW1, 2, 3, 4, 5a, 5b, 5c, 5d, 5e, 6, 7, 8, and 9. MBL has 18 specializations: MBL1, 2a, 2b, 3, 4, 5a, 5b, 5c, 6, 7, 8a, 8b, 8c,9, 10, 11, 12, and 13. These ecosystem associations are elementary models in the sense that they can be uniquely identified based on elevation, aspect, overstory, understory, shrub, herb, and soil data.

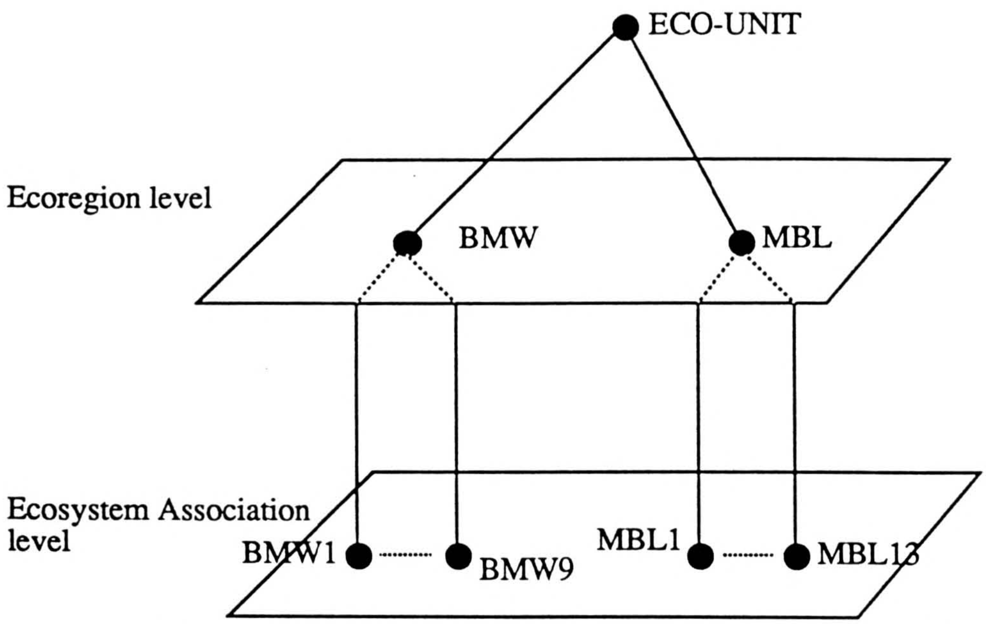


Figure 3: specialization hierarchy for the Duck Mountain, Manitoba area

The primitives in this system are polygons which are the result of overlaying a geographic, forest cover, and soil map. Each polygon is described by a variety of attributes such as area, perimeter, and neighbors. A variety of features can be derived from the maps. These include: elevation, aspect, forest cover, soil texture, drainage, and organic thickness. The mass functions were provided by a forest ecologist with expertise in the local classification system. For the ecosystem domain the mass functions are seen as a subjective expression of the probability of the occurrence of a particular ecoregion / ecosystem association given a particular feature value. These probabilities are based on the ecologist's field research and experience.

Not all input data are necessarily equally reliable. Each feature was therefore assigned a confidence value such that the sum of all confidence values equals 1. Beliefs were expressed on a scale between 0 and 1 where 0 means neutral and 1 very strong belief. Some mass functions were only partially defined which means that symbolic constraints were applied. Belief values were normalized before the application of an update rule. As well, belief values were always distributed over mutually exclusive subsets of ecosystems which implies that the application of Bayes rule and Dempster's rule leads to the same result.

Two different update rules were experimented with: Weighted Averaging and Dempster's rule as proposed in Shafer (1976). Weighted averaging simply updates beliefs by computing the average of the current belief and the belief introduced by a new feature. In this process both the current and new belief are weighted by their confidence values. We assume the features to be independent which is a requirement for the application of Dempster's rule. We chose these two approaches because of their differences in update behavior. Table 1 demonstrates this behavior in the case of 4 features each of which provides beliefs for two ecosystems: A and B. Each feature supports A with belief 0.4 and B with 0.6. The table shows the resulting beliefs of the sequential processing of the features 1 - 4. Whereas weighted averaging leaves the beliefs unchanged, Dempster's rule systematically weakens the belief in A while strengthening the belief in B. We wanted to find out which update rule leads to a better match in follow up field validation trials.

| Feature | Feature constraints | Weighted Averaging | Dempster's rule |
|---|---|---|---|
| 1 | A [0.4] B [0.6] | A [0.4] B [0.6] | A [0.4] B [0.6] |
| 2 | A [0.4] B [0.6] | A [0.4] B [0.6] | A [0.31] B [0.69] |
| 3 | A [0.4] B [0.6] | A [0.4] B [0.6] | A [0.23] B [0.77] |
| 4 | A [0.4] B [0.6] | A [0.4] B [0.6] | A [0.17] B [0.83] |

Table 1: effects of different update rules

The predictive mapping shell has been implemented in Common Lisp Object System (CLOS). The shell is loosely coupled with a Geographic Information System (GIS). Information exchange with the GIS occurs through a file exchange. Data collection and map overlay takes place in the GIS. Information about polygons and features is output into a formatted ASCII file. The mapping shell takes this file as input. The shell, in turn, creates several output files some of which serve as input to the GIS for the creation of ecosystem maps. For the implementation of the predictive ecosystem mapping we used ArcInfo ™ running on a Unix workstation. A more detailed description of the GIS application can be found in Skye (1993). The spatial uncertainty of input data was also considered. A detailed account of the spatial uncertainty issue has been provided by Crain et al., (1993).

## 4. Results and Discussion.

Table 2 illustrates the results of field tests for the three different ecosystem knowledge bases and test areas. In each test area a number of locations were selected where the ecosystem class had been determined by an expert during field observations. Validation results are reported in three different categories; 1+: The percentage of cases in which the expert's choice is also the system's highest ranked choice, 2+: The percentage of cases in which the expert's choice is among the two most highly ranked choices of the system, 3+: The percentage of cases in which the expert's choice is among the three most highly ranked choices of the system.

Table 2 indicates that Weighted Averaging is consistently better than Dempster/Shafer. Two reasons can be given for the difference in performance. The first one is that, apparently, Weighted Averaging better reflects the reasoning process of the domain experts. Another reason is that, because of perceived data reliability, features need to be weighted. While Weighted Averaging has the capability to account for weights on features, Dempster/Shafer does not, although some efforts in this direction have been undertaken by Yen (1986).

| Test Area | No of Points | Category | Weighted Averaging | Dempster/Shafer |
|---|---|---|---|---|
| West Central Alberta | 52 | 1+<br>2+<br>3+ | 63<br>77<br>85 | 38<br>71<br>81 |
| Manitoba Duck Mountain | 53 | 1+<br>2+<br>3+ | 85<br>92<br>98 | 77<br>92<br>92 |
| Manitoba Sandilands | 17 | 1+<br>2+<br>3+ | 94<br>100<br>100 | 82<br>100<br>100 |

Table 2: field test results for three ecosystem knowledge bases

As mentioned in the introduction, this paper had three main objectives:
1. To explain the concept of Evidential Discrimination Reasoning and its implementation as a classification shell.
2. To show how the shell can be used in the design of a decision support system capable of inferring ecosystems from topography, forest, and soil data.
3. To show how such a system can be used as a cost effective tool in the creation of an ecosystem-based land inventory while at the same time improving the quality of the classification process.

The key idea behind Evidential Discrimination Reasoning (EDR) is the organization of the interpretation domain into a specialization hierarchy of models. This enables us to represent the interpretation domain at different levels of refinement. The question of how beliefs are inherited in specialization hierarchies has not been addressed well in the AI literature. We have shown how the traditional notion of a specialization hierarchy can be extended for the purpose of representing and propagating beliefs. Beliefs are represented by mass functions which express the symbolic and numeric constraints imposed by features on models. A feature can constrain a model at any level of specialization. We have adopted the strong inheritance approach in which each node in the hierarchy inherits both the symbolic properties and the mass functions of its parent.

The principle of least commitment allows us to conduct the classification process in an efficient manner. EDR becomes particularly advantageous when the specialization hierarchy is very large (i.e. more than 100 leaf nodes). At any time the current interpretation results are represented at the level of specialization required by the features processed thus far. The order in which features are processed, although irrelevant with respect to the final result, is important with regards to processing efficiency. The classification process will be most efficient if the features that invoke the more general models are processed first. Such an order will both keep the average size of each model listing down, as will it shorten the time to process each feature.

The classification shell will work most efficiently if the structure of the specialization hierarchy closely follows the constraints imposed by the features. In ecosystem classification, for example, one feature (elevation) constrains models at the ecoregion level, whereas all other features constrain models at the ecosystem association level. Given an exhaustive listing of features and the models constrained by them, it is possible to automatically construct specialization hierarchies (Mulder, 1987; Muise, 1987).

As mentioned before, the issue of belief representation can be separated from belief updating. The strong inheritance approach provides a prescription for the distribution of a model's belief over its specializations in the hierarchy. Belief updating, on the other hand, can be treated as a separate issue. The mapping shell allows for the use of different update rules appropriate for the domain of application.

The second objective of this paper was to show how the EDR approach can be applied to the problem of predicting ecosystems from topographic, forest cover, and soil data. The classification shell can solve any classification task that can be represented in terms of primitives, features, and models organized in a specialization hierarchy. The discussion in section 3 has shown, how the ecosystem prediction task can be translated into the EDR format.

The third objective was to show how a decision support system for predicting ecosystems from existing land resource data can be used as a cost effective tool in the creation of ecosystem-based land inventories. An ecosystem-based land inventory can be created in two different ways:

1. By a detailed and systematic field survey effort during which a team of field ecologists map a forest management area into ecosystem associations followed by digitization of the data collected.
2. By "Naiatization" of the area.

The first alternative is both expensive and time consuming. Application of the NAIA system, on the other hand, can save both costs and time. The NAIA system uses data that already exists. In combination with a Geographic Information System it creates an ecosystem map for the area that displays, for each polygon, the ecosystem association with the strongest belief. In addition, the system can show for each polygon the complete list of possible models and belief values. The system can be used as a focal point for the field work. For example, if the predictive map shows a particular ecosystem association as the favorite for a particular area and the belief value for this association is consistently much higher than the value of its closest rival, then a few ground observations may suffice to verify the system's predictions. On the other hand, if the system's predictions are highly ambiguous (i.e. many models with nearly equal belief values), then more ground observations will be necessary. Thus, the system can be used as a basis for prioritizing where field work should be done to maximize the quality of the ecomap for decision making.

Another feature of the system is its capability to bring out inconsistencies between data and models. This happens, when the system fails to classify one or more polygons. Such a failure indicates, that there is a gap or incorrectness in the ecological models for the area, or it may indicate a problem with some of the map data used.

However, the main advantage of the NAIA system remains that it reduces the need for field work while it eliminates the need for digitization as the results are already in digital form. This is where NAIA can save costs. The reduced need for field work also implies that the job can be done by a single forest ecologist rather than by a team. This ecologist can either build the knowledge base, if one is not already available, or he/she can refine the mass functions to accommodate local conditions in the area to be classified. As different ecologists do not necessarily use the exact same classification criteria this will enhance the consistency (and thus, the quality) of the classification process.

## 5. Summary

This paper has described the concept of Evidential Discrimination Reasoning and its implementation as a classification shell. The object-oriented shell recognizes three different types of objects: primitives, features, and models. Models are organized in a specialization hierarchy. The traditional notion of inheritance has been extended to include the representation and propagation of beliefs. The classification process is guided by the principle of least commitment and the principle of graceful degradation. A decision support system for predicting ecosystems from existing land resource data has been successfully implemented in the shell. This system can assist with the creation of ecosystem-based land inventories with the potential of saving costs while at the same time improving the quality of the classification process.

## 6. Acknowledgments

The research described in this paper has been supported by a grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada (OG0121247) to the first author, and by the Canada - Manitoba partnership agreement in Forestry (contract to Hughes Aircraft). The NAIA DSS was developed under a Joint Research Agreement between Hughes Aircraft of Canada Limited, Spatial Data Systems Division, and the Alberta Research Council, with support from the University of Calgary, Department of Geomatics Engineering. The authors are grateful to Dan Jaliff, Marlene Jones, and Breen Liblong for comments on a draft version of this paper. Doug Allan of Forestry Canada (Northern Forestry Centre), Canadian Forest Products (Grande Prairie division), Weldwood of Canada (Hinton division), and Len Knapik of Pedocan Land Evaluation Ltd. are gratefully acknowledged for providing data and expertise. ARCINFO is a Geographic Information System marketed by ESRI.

## 7. References


Brachman R.J., What "ISA" Is and Isn't, in *Proceedings, 4th National Conference of the Canadian Society for Computational Studies of Intelligence,* Saskatoon, Canada, 1982, pp. 212-221.

Brooks R.A., Symbolic Reasoning among 3-D Models and 2-D Images, *Artificial Intelligence* **17**, No 1-3, 1981, 285-348.

Buech R.R., Martin C.J., Rauscher H.M., An Application of Expert Systems for Wildlife: Identifying Forest Stand Prescriptions for Deer, *Applications of AI in Natural Resources, Agriculture, and Environmental Sciences,* Vol. 4, No 1, pp. 3 - 10, 1990.

Crain I.K., Gong P., Chapman M.A., Lam S., Alai J., Hoogstraat M., Implementation Considerations for Uncertainty Management in an Ecologically-oriented GIS, *Proceedings of GIS 93,* Vancouver B.C., 1993, pp. 167-172.

Corns I.G.W., Annas R.M., Field Guide to Forest Ecosystems of West-Central Alberta, Northern Forestry Centre, Canadian Forest Service, Edmonton, 1986.

Hanson A.R. and Riseman E.M., VISIONS: A Computer System for Interpreting Scenes, in *Computer Vision Systems*, (Hanson A.R. and Riseman E.M., Eds.), pp. 303-334, Academic Press, New York, 1978.

Ironside G.R., The Canada Committee on Ecological Land Classification (CCELC): History, Objectives and Activities in Land / Wildlife Integration, Ecological Land Classification Series no 11, Lands Directorate, Environment Canada, 1980.

Jones R.K., Next Generation Forest Site Classification: Ecologically - Oriented Predictive Mapping Technology, *Proceedings of GIS 93,* Vancouver B.C., 1993, pp. 143-152.

Knapik L.J., Russell W.B., Riddell K.M., Stevens N., Forest Ecosystem Classification and Land System Mapping Project, Duck Mountain, Manitoba, Northern Forestry Centre, Canadian Forestry Service, Edmonton, Alberta, July 1988.

Knapik L.J., Ellis R.A., Phillips B., Forest Ecosystem Classification and Land System Mapping Pilot Project, Sandilands, Manitoba, Northern Forestry Centre, Canadian Forestry Service, Edmonton, Alberta, March 1989.

Mackworth A.K., Consistency in Networks of Relations, *Artificial Intelligence* **8**, No 1, 99-118, 1977.

Mackworth A.K., Mulder J.A., Havens W.S., Hierarchical Arc Consistency: Exploiting Structured Domains in Constraint Satisfaction Problems, *Computational Intelligence* **1**, no 3 - 4, 118 - 126, 1985.

Mitchell W.R., Green R.E., Identification and Interpretation of Ecosystems of the Western Kamloops Forest Region, Land Management Handbook no 2, Province of British Columbia, Ministry of Forests, 1981.

Muise J., Constructing Discrimination Graphs to Represent Visual Knowledge, M.Sc. thesis, Department of Mathematics, Statistics, and Computing Science, Dalhousie University, 1987.

Mulder J.A., An Algorithm which Automatically Constructs Discrimination Graphs in a Visual Knowledge Base, *Proceedings of the 10th Int. Joint Conf. on Artificial Intelligence,* Milan, Italy, 1987, pp. 855-857.

Mulder J.A., Discrimination Vision, *Computer Vision, Graphics, and Image Processing* **43**, 313 - 336, 1988.

Pearl J., On Evidential Reasoning in a Hierarchy of Hypotheses, *Artificial Intelligence* **28**, 9 - 15, 1986.

Planchet B., Credibility and Conditioning, *Journal of Theoretical Probability* **2**, 289-299, 1989.

Shafer G., A Mathematical Theory of Evidence, Princeton University Press, 1976.

Shafer G., Perspectives on the Theory and Practice of Belief Functions, *International Journal of Approximate Reasoning* **4**, no 5-6, 323-362, 1990.

Shafer G., Logan R., Implementing Dempster's Rule for Hierarchical Evidence, *Artificial Intelligence* **33**, 271 - 298, 1987.

Skye D., Naia: Ecologically-oriented Spatial and Knowledge-based Framework to Support Forest and Land resource Management, *Proceedings of GIS-93,* Vancouver B.C., 1993, pp. 161-166.

Smets P., About Updating, *Proc. of the 7th AAAI Conf. on Uncertainty in Artificial Intelligence,* Los Angeles, 1991, pp. 378-385.

Yen J., A Reasoning Model Based On An Extended Dempster-Shafer Theory, *Proceedings of the 5th National Conference of the American Association of Artificial Intelligence,* Philadelphia, 1986, pp. 125-131.

**Biographic Information**

Jan Mulder is a research scientist in the Advanced Computing and Engineering Department of the Alberta Research Council in Calgary, Alberta. He received the "kandidaats" and "doctoraal" degrees in experimental psychology from the University of Leiden in 1974 and 1979, respectively, and the M.Sc. and Ph.D. degrees in computer science from the University of British Columbia in 1979 and 1985, respectively. His research interests include: constraint satisfaction, diagnostic reasoning, and knowledge-based computer vision and spatial data systems.

Ian Corns is a research scientist and project leader of the Geographic Information Systems, Forest Inventory and Site project at the Northern Forestry Centre of Forestry Canada in Edmonton, Alberta. He received his training (B.Sc. Botany 1969; M.Sc. Plant Ecology 1972; and Ph.D. Forest Soils 1977) from the University of Alberta. His research interests include: land classification and evaluation, boreal forest ecology, and decision support systems for mixedwood management.

# AI IN AGRICULTURE, NATURAL RESOURCES AND ENVIRONMENTAL SCIENCES

W25

IJCAI 93

INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

CHAMBERY SAVOIE FRANCE

13th INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

IJCAI '93 Workshop W25:

# AI in Agriculture, Natural Resources and Environmental Sciences

13th International Joint Conference on AI

Chambery, Savoie
France

30 August 1993


Organizing Committee

Dr. Nicholas D. Stone (Chair)
Virginia Tech
Blacksburg, Virginia, USA

Dr. Roger Martin-Clouaire
Institut National de la Recherche Agronomique
Castanet-Tolosan Cédex, France

Dr. Richard L. Olson
USDA-ARS, Mississippi State University
Mississippi State, Mississippi, USA

Dr. Hannu Saarenmaa
Finnish Forest Research Institute
Helsinki, Finland

Contents & Schedule:

2:05 **Planning in agriculture and the environment: constraint networks and knowledge-engineering revisited**
S. A. Osinga, M. S. Elzas, and R. J. M. Hartog, Wageningen Agricultural University, The Netherlands

2:30 **BREAK**

2:45 **Object-oriented constraint satisfaction planning for whole-farm management**
Nicholas D. Stone, Virginia Tech, USA

3:10 **Crop management planning as a CSP with advanced features**
Roger Martin-Clouaire and Jean-Pierre Rellier, INRA, Castanet-Tolosan Cedex, France

3:35 **NAIA: A decision support system for predicting ecosystems from existing land resource data**
Jan A. Mulder and Ian G. W. Corns, Alberta Research Council and Forestry Canada

4:00 **BREAK**

4:15 **The application of knowledge-based techniques to support resource management decisions**
David Robertson, Mandy Haggith, Gill Kendon, Jaume Agusti, and David Goldsborough, University of Edinburgh, Scotland, Blanes, Spain, and Van Hall Institute, The Netherlands

4:40 **A framework for multi-agent problem solving and construction of decision support systems for natural resource management through the modeling of an expert enterprise**
Hannu Saarenmaa, Jari Perttunen, Jouni Vdkevd, and J. Michael Power, Finnish Forest Research Institute, Finland, and Forestry Canada

5:05 **Discussion & Summaries**

5:45 **Concluding Remarks & Adjourn**