An expert system for capturing and applying soil survey tacit knowledge to automatically link soils to landform position in soil-landform models

R. A. MacMillan[1], W. W. Pettapiece[2] and J. A. Brierley[2]

ABSTRACT

Soil survey is a paradigm based science which relies heavily on the application of conceptual soil-landscape models based on tacit expert knowledge. The tacit knowledge is generally acquired by systematic field observation and recording of repeating relationships between soils and landform position. A consistent deficiency of many soil survey databases has been the lack of a mechanism for capturing and recording such tacit expert knowledge. A case in point is the recently completed 1:100,000 scale digital soils database prepared for the province of Alberta (AGRASID). This database provided a list of up to 6 soils that were believed to occur within each mapped soil polygon and a landscape code to identify a type landscape considered to best describe the area enclosed by each polygon. It did not include any mechanism for indicating which portion or portions of the landscape each soil was most likely to occur in or for associating each soil with a specific set of landform conditions such as slope gradient or slope length. A procedure was developed to facilitate capture and application of local tacit knowledge in order to associate each soil in each AGRASID polygon with its most likely landform position or positions. The procedure, and associated computer programs, were applied to the AGRASID digital soils database which consisted of over 50,000 individual soil polygons which referenced more than 1,500 different named soil series (or non-soils). The procedures capture and codify the essential components of local expert tacit knowledge required to effectively associate each listed soil in any given polygon with its most likely landform position(s). The concepts on which the procedures were based are simple but effective. They are sufficiently generic as to be applicable to any region in which there is reason to expect consistent and repeating relationships between soils and landform position. They require no special data beyond that which is normally available for any named soil series. The tacit knowledge base, and the procedures used to capture and apply it, may be of as great a potential value as the expanded soil-landform model database created by this specific application of the procedures.

INTRODUCTION

Soil survey as a paradigm based science founded on tacit knowledge

Soil survey has long been acknowledged to be a paradigm based science (Hudson, 1992) in which the underlying hypothesis is that the location and distribution of soils in the landscape is predictable (Arnold, 1979; 1988; Miller et al., 1979). This spatial distribution is widely agreed to be a function of the five soil forming factors of Jenny with topography often playing a dominant role locally. Application of soil survey techniques to produce a map has been described as an iterative exercise involving collecting initial data and observations to support the development of conceptual models of repeating soil landform patterns followed by field testing of the conceptual models (Arnold, 1979, Miller et al., 1979; ECSS, 1987). Once validated by field testing, these conceptual soil-landscape models become the basis for efforts to identify and map the locations of similar landforms with an assumed similar assemblage of soils.

A common misconception is that the sole purpose of soil survey is "to make a soil map" (ECSS, 1987) with the map representing the principal means of capturing and expressing knowledge of soil-landform relationships acquired during the survey process (Hudson, 1992). One significant drawback of many soil surveys has been their failure to explicitly capture tacit expert knowledge

[1] LandMapper Environmental Solutions, Edmonton, AB E-mail: bobmacm@planet.eon.net
[2] Agriculture and Agri-Food Canada, Lethbridge Research Centre E-mail: pettapw@em.agric.gov.ca

on soil landform relationships acquired during the mapping process (Hudson, 1992). At best, soil survey reports may contain a series of notional cross sectional diagrams which attempt to record and portray conceptual models developed during the survey, but there is seldom any attempt to explain why the soils are distributed as they are.

No formal, systematic mechanism has been proposed, or widely adopted, for capturing local tacit expert knowledge regarding why certain soils occur in certain landform positions or in certain common combinations. In the absence of a formal mechanism, the common means of passing on local expert knowledge has been through a prolonged soil survey field apprenticeship under the guidance of an experienced mentor. The combined body of tacit knowledge regarding where certain soils occur in the landscape and why they occur there inevitably becomes stored in the minds and memories of a very limited number of local experts who have passed through this program of field experience and indoctrination. An unavoidable consequence of this approach to storing tacit knowledge is that it tends to disappear with the retirement of each successive generation of soil surveyors. A corollary to this, is that a significant component of local tacit knowledge is inevitably relearned, at considerable expense, by each new generation of trained field personnel. Now that ongoing, publicly funded, soil survey programs have been virtually eliminated in Canada (and in many other locales) the opportunity to continue to pass on expert tacit knowledge through a prolonged field apprenticeship no longer exists for the next generation of local soil experts.

It has been argued that this unrecorded tacit knowledge regarding the patterns of distribution of soils in the landscape has as great a potential value as the soil maps produced through its application (Hudson, 1992). In one Canadian example, it has been reported that an experienced soil surveyor in Saskatchewan (Jim Ellis??) argued against a plan to initiate a new round of more detailed soil surveys in that province in the 1950s (Dan Pennock, personal communication, 1999). His proposal was that the effort should be concentrated instead on building, testing and recording conceptual models describing the most common patterns of distribution of soils by landform position. He proposed that these conceptual models, once validated, should be appended to the existing small scale soils maps to help users appreciate, and predict for themselves, the most likely patterns of distribution in the landscape of individual named soils within the broadly mapped soil associations. This farsighted vision was, unfortunately, not adopted and soil surveyors in Saskatchewan, and throughout Canada, continued to make "one more soil map" until support and funding for these mapping activities was withdrawn.

A successful mechanism for capturing and cataloguing local tacit knowledge on soil landform relationships has significant potential to make existing maps more useful and to guide in the production of a next generation of more detailed maps. If, as suggested by Ellis, one captures the basic understanding of where and why certain soils occur in certain locations in the landscape, it may be relatively easy to produce more detailed, site level maps, starting with widely available, generalized soil maps.

The Alberta situation

The situation in 1999 regarding soil survey information in Alberta, Canada offers a classic example of the need and opportunity to capture and apply local tacit knowledge in a formal systematic way.

In 1993 Alberta began a project whose aim was to collate all existing soil survey maps and information into a single, seamless, digital soil survey database covering the entire agricultural portion of the province (the White Area) (CAESA, 1993). This project was completed and a final digital database was available for distribution in 1998 (AGRASID reference). The 1:100,000 scale AGRASID digital soils database covered an area of over 1.7 million km^2 and consisted of over 50,000 individual soil polygons and over 1,500 named and described soils.

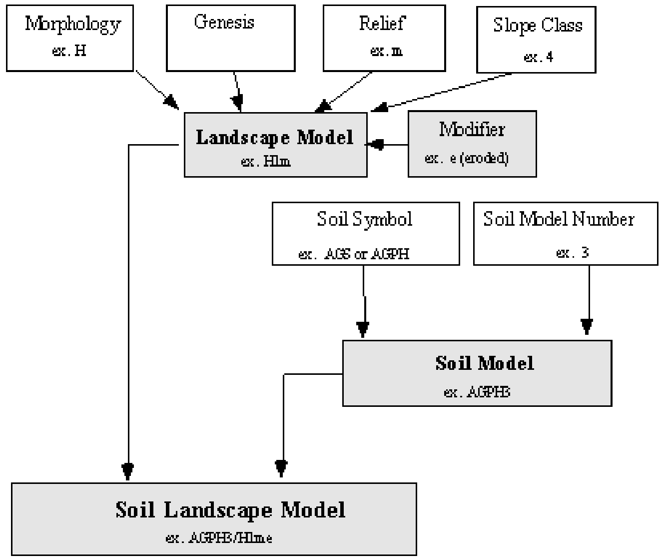

Figure 1 Components of an AGRASID Soil Landscape Model (source: Nikiforuk, 1997 Figure 2.3)

Documentation prepared for the AGRASID database (Nikiforuk, 1997) takes great pains to describe procedures that are claimed to result in the production of a soil-landscape model for each unique combination of soils and landforms in the AGRASID database (Figure 1).

In fact, the AGRASID entity labeled as a soil-landscape model falls significantly short of satisfying all of the requirements of a complete soil-landscape model. As presented, it essentially represents a collation of two separate and non-related symbols, one representing a soil model and the other a landscape model.

The landscape model symbol identifies a code for one of 64 different types of landform defined for use in the AGRASID project. Each landform type is described very generally in terms of its surface form, its assumed geological origin (genesis), its dominant relief and its dominant slope gradient class. Code descriptions are recorded in a table in the AGRASID on-line documentation, but users must be reasonably familiar with their meanings and definitions in order to interpret the symbol while viewing the master soils database.

The soil model reflects a convention widely used in Alberta prior to AGRASID in which specific combinations of dominant, co-dominant and significant soils are identified by means of an alpha-numeric symbol. The alpha portion of the symbol is based on utilization of either a 3 letter alpha code to identify a single dominant soil or a 4 letter alpha code based on concatenation of the first 2 letters of the 3 letter soil codes for 2 co-dominant soils. The number portion of the symbol is used to identify a specific combination of dominant and significant soils (e.g. Chernozemic with significant amounts of Solonetzic). This "soil model" represents an attempt by the soil map compilers to recognize and map repeating sequences of soils that commonly occur in association with one another.

What is conspicuously absent from the AGRASID soil-landscape model is some mechanism for informing the user about the distribution of listed soils within the main landform positions of the listed landscape type. One is presented with 2 separate pieces of information, one identifying the landscape as having a particular set of characteristics and another identifying a list of the main soils believed to occur in the landscape. There is no capability to inform the user as to which of the listed soils are most likely to occur in the each of the main landform positions within the landscape of interest.

It is almost certain that the individual responsible for compiling the digital soils information possessed a working soil-landscape model based on expert tacit knowledge and used this model to assist in assigning a description to the AGRASID polygon. It is equally likely that this individual, and others with a comparable level of local tacit expert knowledge, would be able to review the list of named soils and to assign each soil to its most likely position or positions in the landscape. For a number of reasons, including concerns about costs and reliability, this tacit knowledge concerning where in the landscape each soil was most likely to occur, was not captured or recorded in the AGRASID digital database.

This lack of ability to associate each named soil in an AGRASID polygon with one or more positions in the landscape where it is most likely to occur is a significant impediment to full and effective use of the AGRASID database. Many models, both deterministic and logic based, require an ability to associate a soil, and perhaps more importantly the properties of a soil, with a specific portion of the landscape characterized by a defined slope gradient and slope length. This capability is lacking in the initial AGRASID digital database, and indeed is generally not widely available in most soil maps produced in Canada.

An additional feature of the AGRASID digital database is that it represents a final legacy of soil survey activity in Alberta. Compilation of this seamless, comprehensive digital database was the final activity undertaken by publicly funded soil survey organizations in Alberta. By the time the AGRASID database was completed, both federally and provincially supported soil survey organizations had been permanently shut down and their experienced staff were either retired or employed in the private sector as consultants. With these retirements and reassignments, a considerable portion of local expert tacit knowledge regarding relationships between soils and landform position was effectively being lost. This raised the question of whether some mechanism might not be found to capture, retain and apply this valuable, but threatened, knowledge base. It was considered unlikely that any future generations of soil surveyors would be afforded the opportunity of acquiring comparable levels of local tacit knowledge through leisurely and mentor-guided acquisition of relevant field experience.

OBJECTIVES

The general objective of the present project was to develop, apply and evaluate a generic set of procedures for capturing local tacit knowledge to assist with automated allocation of soils to their most likely positions in the landscape.

The specific objectives particular to the 1:100,000 scale AGRASID digital soils database were to:

- Develop a comprehensive quantitative database describing the major morphometric characteristics of each of the 64 type landforms defined for use in the AGRASID project.
- Devise and apply a systematic set of procedures for automatically associating every soil listed as occurring in an AGRASID polygon with its most likely location in the landscape.
- Devise and apply a systematic set of procedures for automatically constructing comprehensive soil-landform models for every unique combination of soils and landform in the AGRASID digital database.
- Evaluate the utility and reliability of the soil-landform models produced by application of the procedures.

METHODS

The procedure used to automatically allocate any number of soils listed as occurring in a given map polygon in the AGRASID digital soils database to their most likely landform position (or positions) involved the following main steps.

- Creation of a database describing each of the 64 landform types defined for use in the AGRASID digital mapping program and specifically providing an estimate for the percent extent of 4 simple conceptual landform classes (UP, MID, LOW, DEP) for each landform type.
- Consideration of the basic data set of information recorded for every named soil in the NSDB Soil Names File (SNF) in order to identify and select components of this basic data set that were known, or at least believed, to reflect relatively consistent and predictable patterns with respect to location of the described attribute in the landscape.
- Creation of a rule base of expert knowledge and opinion regarding the relative likelihood of each of the ***n*** classes defined for each of the selected soil attributes occurring in each of the 4 defined landform positions.
- Creation and application of a program to convert the expert judgement and belief captured in the rule base(s) into numbers scaled from 0-100 which reflected the relative likelihood of each soil in the SNF occurring in each of the 4 defined landform positions (UP, MID, LOW, DEP).
- Creation and application of a program to read in a list of named soils and a landform type code for each unique soil polygon in the AGRASID digital data base and to identify which of the listed soils were deemed most likely to occupy each of the 4 defined landform classes.

Creation of a landform model database (SLM_LMD)

The AGRASID digital soils database contained an alpha-numeric code which identified which of 64 conceptual landform types defined specifically for the AGRASID project was considered by the map compiler to best describe the shape and pattern of the landforms observed within a given map polygon.

Unfortunately, the 64 conceptual landform types were described in only a very general way in the AGRASID procedures manual and documentation (ref here). These descriptions consisted of a rough estimate of the range in relief that was typical for each landform type and a listing of the slope class or slope classes that were considered to best describe the landform. There was no explicit documentation of whether the listed slope classes represented the classes occupying the largest proportion of the mapped polygon (dominant) or whether they represented limiting or controlling slope classes that occupied an important (but not dominant) proportion of a mapped polygon. A key requirement for the present project, which was not provided in the AGRASID data base, was an estimate of the extent of each type landform that could be considered to belong to each of the 4 simple conceptual landform positions on which the present allocation procedures were designed to act.

Since an estimate was required for the proportion of each of the 4 landform positions in each of the 64 type landscapes and since it was not already provided by the AGRASID documentation, it was necessary to first construct a database containing the required data. Detailed, site level (section or quarter section) digital elevation model (DEM) data sets were available for about 16 sites that were each considered to be representative of at least one of the 64 defined AGRASID landform types. These16 sites represented the most common AGRASID landform types which collectively accounted for more than 90% of the total area mapped for AGRASID. The 16 sites were processed to define the location and extent of 4 simple landform elements using a recently developed procedure for automatically segmenting landforms into landform elements (MacMillan et al., 1998, 1999).

The location and percent extent of each of the 4 simple landform positions (UP, MID, LOW, DEP) were computed for each of the 16 sites for which DEM data were available. The computed proportions were then rounded to the nearest 5% and entered into a new landform model database (SLM_LMD). Data for the 16 sites with computed distributions were then used to assist in manually estimating appropriate values for the distribution of the 4 simple landform position classes for the remaining 48 landform types for which detailed DEM data were not currently available. The 16 sites represented landform types occupying the full spectrum of scales of relief, slope gradient, slope length and drainage complexity commonly found in agricultural landscapes in Alberta. They were found to be highly useful for assisting in manually estimating and assigning appropriate values to the remaining landform types.

The DEM data for the 16 representative sites were also analyzed to compute and record statistics for a number of quantitative descriptors of landform morphology. Both cumulative and classed frequency distributions were computed for slope gradient, aspect, profile and plan curvature, several measures of local relief, several measures of slope length and several measures of watershed size, density and degree of integration of surface water flow at each site (MacMillan and Pettapiece, 199x). The quantitative descriptions were further enhanced by computing means and frequency distributions for several measures of slope gradient and slope length for each of the 4 simple landform classes at each site. These data were not required for the procedures used to allocate soils to landform position, but they were viewed as necessary for effective use of the resulting soil-landform models.

Identification of soil attributes related to landform position

The procedure devised to allocate soils to their most likely landform position(s) was based on a fundamental assumption that reliable and consistent relationships could be defined between landform position (as represented by the 4 defined landform classes) and a limited number of soil attributes or classifications. An alternative approach of attempting to capture expert knowledge and opinion regarding the patterns of spatial organization of every known combination of mapped soils (the named map units in the AGRASID data base) was considered too complex, demanding and prone to error to be successfully realized.

The intent was therefore, to construct a knowledge base in which selected soil attributes or classifications could be interpreted and assigned a relative likelihood of occurring in each of the 4 defined landform classes. The knowledge base was to be founded on expert knowledge and opinion regarding local relationships between landform position and these selected soil attributes. It was necessary that the procedure use only soil attributes and classifications for which data were consistently available for all soils to be considered. This effectively limited the list of available characteristics to those stored in the standard National Soils Data Base (NSDB) soil names file (SNF) and soil layer file (SLF).

Creation of an expert rule base relating soil attributes to landform position

The process used to capture and codify expert judgement and beliefs on how each of the 6 soil attributes was related to landform position was remarkably simple and straightforward.

The knowledge was expressed as an integer number between 0 and 100 which was meant to express the degree to which each class of each of the 6 soil attributes was believed to be likely to occur in each of the 4 defined landform positions (see Figure 2).

Final values were arrived at following a series of discussions and arguments aimed at achieving consensus among the three authors regarding the degree to which each class of each of the selected attributes was likely to occur in each of the 4 defined landform positions. The authors possessed a combined total of over 80 years of field experience classifying, mapping and correlating soils in Alberta and western Canada. Four different rule bases were created initially, one for each of the major eco-regions in the agricultural portion of Alberta corresponding to the Brown, Dark Brown, Black and Dark Gray soil zones..

Computing likelihood of a soil occurring in each of 4 landform positions

The overall likelihood of each soil in the Alberta SNF occurring in each of the 4 defined landform positions was computed as a simple weighted average of the likelihood values assigned to each of the 6 selected soil attributes for each soil (see Table 1).

Table 1 Illustration of calculation of a weighted average for overall likelihood of occurrence in the 4 landform positions for 3 example soils

Soil Attribute	Drainage	Salinity	Calc	Parent Material	Variant	Sub Group	Weighted Mean
Attribute Weight	2	1	1	1.5	1.5	3	
Balzac (BZC)	P	S	S	L14		R.HG	
UP	5	5	80	30	0	5	18.2
MID	20	30	40	50	0	10	25.3
LOW	70	90	60	90	0	55	69.4
DEP	95	100	75	95	0	95	93.2
Angus Ridge (AGS)	W	N	M	M4		E.BL	
UP	85	90	70	90	0	75	81.2
MID	90	60	60	80	0	93	82.2
LOW	60	10	50	40	0	70	52.9
DEP	10	20	65	20	0	10	19.4
Eroded Angus Ridge (AGSer)	W	N	M	M4	ER	E.BL	
UP	85	90	70	90	100	75	84.0
MID	90	60	60	80	55	93	78.2
LOW	60	10	50	40	10	70	46.5
DEP	10	20	65	20	1	10	16.7

The likelihood values for each of the 6 individual soil attributes represented a quantitative expression of the combined beliefs of the 3 authors regarding how likely each class characteristic of a given soil was to occur in each landform position (see Table 1). Expert judgement was also used to assign different weights to each of the 6 selected attributes in order to recognize some as more reliable and significant indicators of landform position than others.

A program (SLMSTEP1) was written in the xBase database programming language to read in and apply the appropriate rule base in order to compute the overall likelihood of each soil in a given eco-region occurring in each of the 4 landform positions. The program simply determined the class codes for a given soil for each of the 6 selected attributes as listed in the SNF data base (the alpha-numeric codes listed for each soil in Table 1). It then accessed the appropriate database table containing the rules which expressed expert judgement about how likely a soil with a given class value for a given attribute would be to occur in each landform position.

The likelihood values associated with the current class code of the current attribute were recorded for likelihood of occurring in each of the 4 landform positions (UP, MID, LOW, DEP). Each landform position was considered in turn and the likelihood value for each of the 6 attributes was multiplied by the weighting factor associated with that attribute (Attribute Weight in Table 1). Each of the 6 attributes was likewise considered in turn and a sum of the individual weighted products was maintained for each landform position. Any attribute with a value of 0 was treated as a missing value and was not included in calculation of the weighted average (see Variant for BZC and AGS in Table 1). This was required to deal with cases where a particular soil had no variant associated with it. In such cases, it is impossible to assign a meaningful value to variant and so the variant attribute should not be included in calculation of the overall likelihood score. Once all 6 attributes were considered and the sum of their weighted products determined, the total sum was divided by the total sum of the weights of only those attributes actually included in the calculation (e.g. weights of attributes assigned a missing value identifier of 0 were not included in the total weight).

Table 2 Structure and content of the additional fields added to the modified SNF file (see Table 1)

Field Name	Field Type	Field Length	Field Dec	Description
SLM_ZONE	N	1	0	Re-code of SCA numbers into 1 of 4 zones for which different rule basis were created
UP	N	4	1	Value computed for likelihood of occurring in an upper landform position (UP)
MID	N	4	1	Value computed for likelihood of occurring in an mid landform position (MID)
LOW	N	4	1	Value computed for likelihood of occurring in an lower landform position (LOW)
DEP	N	4	1	Value computed for likelihood of occurring in an depressional landform position (DEP)

The final weighted average for a given soil for likelihood of occurring in each of the 4 landform positions was recorded in a modified version of the original Soil Names File. The modified SNF had 5 additional fields added (Table 3). The field SLM_ZONE was inserted to assign all soils in the SNF to 1 of the 4 zones or eco-regions for which different rule bases had been created. The program used the zone identifier to determine which of the 4 defined rule bases to use when computing likelihood values for a given soil. The other 4 fields were simply locations in which to store the results of the calculation of likelihood of occurring in each of the 4 defined landform positions for each soil in the SNF.

Creation of a data base linking all soils in a polygon to landscape positions

The final step in the process of automatically allocating soils to landform positions was to create and apply a program to link any number of soils, as listed for an AGRASID soil map polygon, to the most likely location (or locations) in the landscape for the specific landform type associated with the polygon of current interest in the AGRASID soil data base.

The program (SLMSTEP2) required as input only the following information :

- a list of unique SNF soil codes and extent codes for all soils present in the polygon of interest (extracted from the AGRASID master soil (MAS) file)
- a unique code identifying one of the 64 types of landform defined for AGRASID (extracted from the AGRASID soil landscape (SL) file)
- access to the landform model database (SLM_LMD) in which were stored the computed estimates of the relative extent of each of the 4 landform positions for each landform type
- access to the modified SNF (SLM_SNF) file in which were stored the computed values for likelihood of occurring in each of the 4 landform positions for every named soil in Alberta.

The program produced as output a new, revised and expanded soil landform model database (SLM_SLM) in which each of the 4 simple landform positions was associated with one or more defined soils, these being the soils with the highest value for likelihood of occurring in each landform position.

The program (SLMSTEP2) considered which soil or soils were most likely to occupy a particular landform position in a specific sequence. Upper landform positions (UP) were considered first, followed by depressional (DEP) then midslope (MID) and finally lower slope (LOW) (Figure 2).

The first step was to determine the extent of the current landform identified as belonging to an upper landform class (UP). Then the soil with the highest likelihood of occurring in an upper landform position (as recorded in the modified SLM_SNF) was identified (soil A in Table 3). If the percent extent of this first soil, as computed from the AGRASID database, was greater than the percent extent assigned to the upper landform class for the current landform type, then the entire percent extent of the upper landform class was associated with this initial soil (A). An entry was made in a temporary working file which indicated that the upper landform position for the current polygon consisted of X% of the polygon and that the entire X% was occupied by soil A.

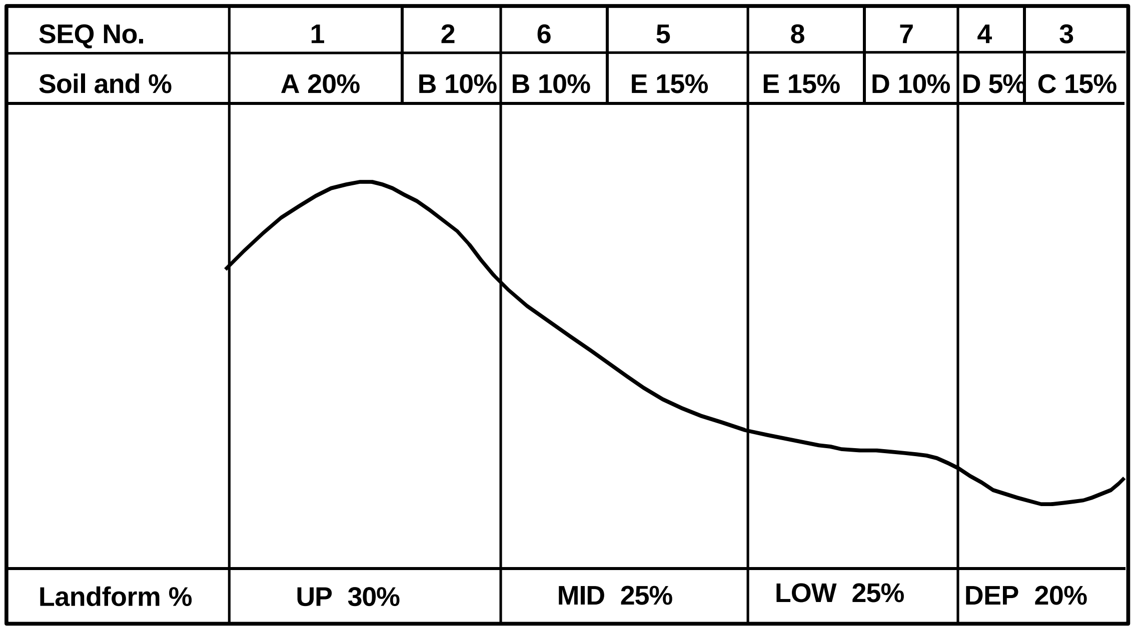

Figure 2 Illustration of the logical sequence used to allocated soils to landform position

A running account was maintained of the total amount of each soil (e.g. soil A) allocated to any landform position. The running total of a given soil allocated to the upper landform position was subtracted from the initial total extent of the soil before allocation to determine how much, if any, of the current soil remained to be allocated. If the recorded percent of soil A was less than the percent of the landscape defined for the upper landform class (UP) then all of soil A was associated with the upper landform element and the running total for the allocated extent of soil A exactly equaled the initial total known extent of soil A. Since soil A had been entirely allocated, and since the upper landform element was not yet fully associated with named soils, it was necessary, in this case, to identify and begin to allocate, the soil with the next highest likelihood of occurring in the upper landform position. This next soil (e.g. B) was identified and if its known initial extent did not exceed the extent of the upper landform element still requiring an associated soil, then the entire percent extent of soil B was associated with the upper landform position and a second entry was made in the temporary working file indicating that X% of the UP landform position was occupied by soil B. This process continued until all of the reported extent of the upper landform position was associated with one or more named soils.

Table 3 Hypothetical data set illustrating the data required to allocate soils to landform position

		Likelihood of Occurring in Each of the 4 Landform Positions			
LF	Percent Extent	30%	25%	25%	20%
Soil	Percent Extent	UP	MID	LOW	DEP
A	20%	87	73	48	12
B	20%	82	76	52	16
C	15%	14	36	81	92
D	15%	16	41	86	78
E	30%	78	83	71	53

The second step involved applying the same procedures as above to allocate the most likely soils to the depressional landform position (DEP). Again, the soil with the highest assigned likelihood of occurring in a depression was first identified and associated with the depressional landform class. If this soil (C) was greater in extent than the reported extent of the depressional landform class for the current landform type, then the entire extent of the DEP class was associated with soil C. If the extent of soil C was less than the reported extent of the DEP landform class then all of soil C was associated with the DEP class and the next most likely soil to occur in a depression was identified (e.g. D). This process was continued until all of the reported extent of the DEP landform class was associated with one or more of the named soils listed for the current polygon.

It was imperative to allocate soils first to the upper then to the depressional landform classes. This was done to avoid situations in which a soil whose likelihood of occurring in a depression was higher than any other listed soil might be incorrectly allocated first to a mid or lower landform position, if these elements were considered before the depressional elements. It was not uncommon for a soil with the highest likelihood of occurring in a depression to also display the largest value for occurring in a lower or mid slope, especially among those soils remaining after the most well drained soils had been allocated to the upper landform positions. For the logic to work properly, it proved necessary to address allocation of soils to the two extremes in the landscape (upper and dep) first and to only then consider the intermediate landform positions.

Having associated the most likely soils with the upper and depressional landform positions, the program proceeded to apply a similar logic to associate the remaining, presently unallocated, soils with first the mid and then the lower slope positions. The process was deemed complete when the entire extent of all 4 slope classes was associated with one or more of the soils listed for the current polygon and the entire extent of all listed soils had also been associated with one or more landform positions.

The temporary working file containing the data for a single AGRASID polygon was sorted and indexed to reorder the soils in a logical manner from highest to lowest landform position and from most likely to least likely soil within each landform class. This resorted data was written back into a permanent database table (SLM_SLM) in which each of the 4 landform positions was associated with one or more of the named soils listed for the polygon. This data table listed the extent of each named soil associated with each of the 4 landform positions as well as the total extent of each landform position class for each polygon of interest.

Each record in the database also included values for the mean and controlling slope gradient and controlling slope length associated with each combination of soil code and landform position. These data were not part of the requirements for allocating soils to landform position, but were included because they were deemed essential for effective use of the proposed soil landform model database.

RESULTS

The landform model database (SLM_LMD)

The complete landform model database created to support the landform allocation procedures is too large to be reproduced here but is listed in Appendix 1.

A portion of the larger database illustrating only the information required to support the automated procedures for allocating soils to landform position is illustrated in Figure 3.

Slm_lmd

Lf_code	Lf_desc	Ups_prop	Low_prop	Mid_prop	Dep_prop
FP1	meander floodplain	10	40	40	10
FP2	braided channel	0	20	50	30
FP3	confined, terraced	10	60	20	10
L1	level plain	0	45	45	10
L2	closed basin	10	40	10	40
L3	level, terraced (not in valley)	15	20	60	5
U1l	undulating - low	20	15	50	15
U1h	undulating - high	25	20	45	10
IUl	inclined & undulating - low	20	20	55	5
IUh	inclined & undulating -high	20	25	50	5
H1l	hummocky - low	30	20	40	10
H1m	hummocky -med	30	25	35	10
H1h	hummocky -high	35	25	30	10
H5l	hummocky over BR - low	30	20	45	5
H5m	hummocky over BR -med	30	25	40	5
H5h	hummocky over BR -high	35	25	35	5
R2l	ridged - low	20	20	55	5
R2m	ridged - med	20	15	60	5
R2h	ridged - high	15	15	65	5
D1l	longitudinal dune - low	20	20	55	5
D1m	longitudinaldune - med	20	15	60	5
D1h	longitudinal dune - high	15	15	65	5
D2l	parabolic dune - low	20	15	45	20
D2m	parabolic dune - med	20	10	50	20
D2h	parabolic dune - high	15	10	55	20
M1l	rolling - low	25	25	45	5
M1m	rolling - med	25	20	50	5

Figure 3 Illustration of a portion of the landform model database (LMD) required to support allocation of soils to landform position

The landform model database (SLM_LMD) contains an estimate of the extent of each of the 4 simple landform elements within each of the 64 landform types defined for the AGRASID database. This is the only information in the LMD required for allocating soils to landform elements. Other information contained in this file was included primarily to support other anticipated requirements such as the application of erosion or runoff models or partitioning the landscape into functional components in order to scale up observations from site studies (e.g. carbon sequestration) using regional databases (e.g. AGRASID).

Soil attributes in the SBF considered to be related to landform position

The NSDB soil names file (Table 4) was reviewed to assess which of its fields of information contained data on soil classifications or attributes that might be interpretable in terms of a consistent pattern of distribution by landform position. It was decided that classifications pertaining to drainage, salinity, calcareousness (CALCAR), soil variants (VARIANT), type and texture of the soil parent material (MAS-PM) and soil SubGroup (SG) were all amenable to interpretation in terms of their relative location in the landscape.

Drainage class was the most obvious and easiest attribute to interpret in terms of landform position. Most landscapes in Alberta demonstrate drainage toposequences in which rapidly to well drained soils predominate in upper landform positions, very poorly to poorly drained soils occupy depressions and moderately well to imperfectly drained soils are found in mid to lower slopes. The distribution of saline and calcareous soils in the landscape is somewhat more problematic, but a commonly held model envisages salinity being more common in depressions and lower slope positions and becoming less common progressing to upslope positions. Calcareous soils were assumed to be relatively common on eroded upper crests and also somewhat more common in lower, toe slope, landform positions.

Table 4 Structure and content of the standard NSDB Soil Names File (SNF)

Field Name	Field Type	Field Length	Field Dec	Description
NEW_SYMBOL	C	7	0	Concatenation of SERIES CODE + VARIANT into a unique code
SERIES	C	24	0	The full name for the Soil Series
VARIANT	C	4	0	Alpha codes for minor variations from the standard series definition
LU	C	1	0	A code to identify the land use under which the soil was described
SCA	N	2	0	A number to identify the soil correlation area in which the soil occurs
DRAINAGE	C	2	0	Alpha codes for the CanSIS defined drainage class of the soil
CALCAR	C	4	0	Alpha codes for the CanSIS defined calcareousness class of the soil
SALINITY	C	4	0	Alpha codes for the CanSIS defined salinity class of the soil
PM1_TEX	C	4	0	Alpha codes for the CanSIS texture class of the upper parent material
PM1_TYP	C	4	0	Alpha codes for the CanSIS class for type of upper parent material
PM2_TEX	C	4	0	Alpha codes for the CanSIS texture class of the lower parent material
PM2_TYP	C	4	0	Alpha codes for the CanSIS class for type of lower parent material
MAS_PM	C	4	0	Unique to Alberta AGRASID single code for parent material type and texture
REPORT	C	20	0	Text listing of the name of the report in which the soil was documented
ORDER	C	2	0	Alpha codes for the CanSIS symbol for Soil Order in which the soil belongs
S_GROUP	C	4	0	Alpha codes for the CanSIS symbol for Sub Group in which the soil belongs
G_GROUP	C	3	0	Alpha codes for the CanSIS symbol for Great Group in which the soil belongs
SG	C	8	0	Concatenated symbol to uniquely classify the soil to the SubGroup level
SG_MOD	C	8	0	Alpha codes for the CanSIS symbol for Sub Group modifier
CORRNOTE	C	254	0	A comments field for recording correlation notes about the soil

Most soils in the Soil Names File do not have a variant associated with them and many of the recognized soil variants were not readily amenable to interpretation in terms of landform position. However, a significant number of soil variant characteristics were considered to exhibit some tendency to occur preferentially in one or more of the defined landform positions. For example, eroded, stony or coarse variants were considered more likely to occur in upper landform positions while saline, gleyed or peaty variants were expected to be more common in depressions and lower landform positions.

There was considerable concern regarding whether parent material could be interpreted in terms of landform position in a manner that was consistent enough to warrant inclusion in the rule base. In an earlier application of the automated allocation procedures (MacMillan, 1997; MacMillan and Pettapiece, 1997b), parent material was first included then removed from consideration. Ultimately, it proved necessary to consider potential relationships between parent material and landform position. This was required because of the rather frequent occurrence of situations in which two or more soils listed as occurring in a given map polygon were identical in all respects except for type or texture of the parent material. A common example would be landscapes characterized by the presence of a veneer or blanket of medium to moderately fine textured lacustrine or aeolian materials overlying till. In these landscapes, soils with a thin veneer over till are most frequently noted in lower landscape positions when they occur in association with till soils with no veneer but are more common in upper landform positions when the associated soils are developed in a thick lacustrine or aeolian blanket overlying till. These, and other, situations dictated that consideration of relationships between topographic position and parent material be included in the analysis.

Classification of the soil at the SubGroup level was the final attribute considered. It was sometimes difficult to decide which landform position or positions a particular soil SubGroup was most likely to occur in but, in most cases, it was possible to define a continuum of likelihood values which captured local expert knowledge regarding which landform positions each SubGroup was most (or least) likely to occupy. A full analysis of all of the logic associated with the association of soil SubGroup classifications with landform position is not attempted here.

It was felt that the overall likelihood of a soil occurring in each of the 4 defined landform positions could be established by computing a weighted average that reflected the combined likelihood of each of the above 6 soil classifications occurring in each of the defined landform positions.

The expert rule base(s) relating soil attributes to landform position

Four separate knowledge basis were constructed for Alberta, one for each of the main eco-regions in the province corresponding to the Brown, Dark Brown, Black and Dark Gray Soil Zones. This was initially done to accommodate expected differences in the spatial association of soil attributes and classifications with landform position under different ecological conditions It was anticipated that the same class of a given soil attribute might exhibit a different preferred location in, for example, the moist, cool Dark Gray soil zone than in the dry, warm Brown soil zone. A portion of the rule base created for the Black Soil Zone is illustrated in Figure 4.

Bl_rules

Attr_no	Attr_ord	Attribute	Class	Up	Mid	Low	Dep
1	0.00	DRAINAGE	NA	0	0	0	0
1	0.01	DRAINAGE		0	0	0	0
1	1.00	DRAINAGE	VR	100	80	20	1
1	2.00	DRAINAGE	R	95	80	35	5
1	3.00	DRAINAGE	W	85	90	60	10
1	4.00	DRAINAGE	MW	50	65	80	40
1	5.00	DRAINAGE	I	20	30	85	70
1	6.00	DRAINAGE	P	5	20	70	95
1	7.00	DRAINAGE	VP	1	10	25	100
2	0.00	SALINITY		90	60	10	20
2	0.01	SALINITY	"-"	90	60	10	20
2	1.00	SALINITY	N	90	60	10	20
2	2.00	SALINITY	W	40	50	70	60
2	3.00	SALINITY	M	20	40	80	90
2	4.00	SALINITY	S	5	30	90	100
3	0.00	CALCAR		30	100	40	90
3	0.01	CALCAR	"-"	30	100	40	90
3	1.00	CALCAR	N	30	100	40	90
3	2.00	CALCAR	W	60	40	45	50
3	3.00	CALCAR	M	70	60	50	65
3	4.00	CALCAR	S	80	40	60	75
3	5.00	CALCAR	V	90	20	80	85
3	6.00	CALCAR	E	100	10	90	95
4	0.00	VARIANT		0	0	0	0

Figure 4 Illustration of the rule base used to capture and quantify expert tacit knowledge

Each unique class of each attribute in the knowledge base (Figure 4) was assigned a value corresponding to its assumed likelihood of occurring in each of the 4 defined landform positions. For some attributes, such as drainage class, little argument occurred and little discussion was required to arrive at a consensus on appropriate likelihood values to assign to each class. In such cases, it was obvious that there was general agreement on conceptual models held by experienced field personnel with regard to how drainage classes were most likely to arrange themselves with respect to landform position.

More discussion was found necessary before some of the more problematic attributes could be assigned likelihood values that could be agreed to by all. In such cases, it was found useful to select an interim value and run the programs to compute overall likelihood of occupying each landform position. Once this was done, the results were examined to assess whether the assigned values produced expected and acceptable results. If the results were not acceptable the assigned values were revisited and altered as appropriate. Surprisingly few iterations were required to identify and address all suspect likelihood values. Relatively few results from the first run were suspect and most were addressed and cleared up by the second run. A third iteration removed all but the most intractable errors.

The creation of 4 separate rule bases was designed to permit an attribute, for example salinity, to be assigned a high likelihood of occurring in a lower landscape position in a one environment and a low likelihood in another. We were concerned with the possibility of reversals in the order of soils or soil attributes in a toposequence under different ecological conditions. For example, in moist environments, the expected sequence of solonetzic soils, from upper to lower landform positions, might be non-solonetzic, then Solonetzic Blacks, then Solods, then Solodized Solonetz then Solonetz and finally Solonetzic Gleysols. In some dryer environments, a sequence that is almost exactly the inverse of this has been reported, with soils classified as Solonetz occurring in the uppermost landform position and Solods or Solonetzic Browns found in the depressions.

The program for computing likelihood of each soil in the SNF occurring in each of 4 landform positions (SLMSTEP1)

The program LSMSTEP1 was run 4 times, once for each of the 4 rule bases created for each of the 4 defined ecoregions, to populate the modified Soil Names File (SNF) with estimates for the overall likelihood of each of its 1570 named soils (or non-soils) occurring in each of the 4 defined landform positions, (see Figure 5).

Hem_snf

New_symbol	Sca	Slm_zone	Drainage	Salinity	Calcar	Variant	Mas_pm	Sg	Up	Mid	Low	Dep
ABC	21	4	MW	N	W		M4	O.GL	77.6	69.4	46.5	26.2
ABCst	21	4	W	N	W	ST	M4	O.GL	84.5	71.5	37.0	13.8
ACE	7	3	W	N	M		L3	O.BL	77.6	78.5	50.3	23.8
ACV	1	1	W	N	M		F2	CA.B	67.9	50.3	43.2	28.2
ADM	12	4	W	N	M		M2	E.BL	67.1	74.1	52.9	19.4
ADMgl	12	4	I	N	M	GL	M2	GLE.BL	37.3	57.0	74.0	42.0
ADY	6	3	W	N	S		M4	O.BL	85.9	80.6	47.1	22.4
ADYgl	6	3	W	N	S	GL	M4	GL.BL	57.3	58.0	64.0	37.0
ADYsa	6	3	I	W	S	SA	M4	O.BL	55.2	57.0	64.5	47.0
ADYxp	6	3	W	N	S	XP	L6	O.BL	86.1	78.3	46.8	19.8
AGH	18	4	I	W	W		F4	GLD.GL	43.8	49.4	73.1	43.5
AGS	10	3	W	N	M		M4	E.BL	81.2	82.2	52.9	19.4
AGSer	10	3	W	N	M	ER	M4	E.BL	84.0	78.2	46.5	16.7
AGSsa	10	3	MW	M	M	SA	M4	E.BL	55.2	64.4	69.5	41.5
AGSsc	10	3	MW	M	M	SC	M4	E.BL	56.5	70.4	68.0	38.5
AGSst	10	3	W	N	M	ST	M4	E.BL	82.5	77.4	49.5	18.0
ALC	17	4	MW	N	W		F4	O.GL	76.8	69.4	46.5	26.2
ALG	20	4	P	N	N		F1	O.LG	27.6	50.0	68.2	80.3
ALGxt	20	4	P	N	M	XT	L14	O.LG	41.0	49.0	63.5	67.3
ALT	4	2	W	N	M		M4	R.DB	88.2	70.6	31.8	15.7
ALTsc	4	2	W	W	M	SC	M4	R.DB	71.5	66.5	45.0	26.4
AMK	21	4	R	N	N		C2	E.EB	82.4	73.5	33.5	20.1

Figure 5 Illustration of the revised Soil Names File (SNF) with all soils assigned values for likelihood of occurring in each of the 4 landform positions

The likelihood values were considered to be new attributes of each named soil (or non-soil) listed in the SNF. They represented a convex combination of the expert beliefs regarding how likely the 6 individual soil attribute classes of a given soil were to occur in each of the 4 landform positions.

It might have been possible to attempt to use expert knowledge and opinion to directly assign each unique soil a relative likelihood of occurring in each of the 4 landform positions without resorting to an intermediate analysis step based on consideration of the 6 selected soil attributes. This would have required experts to consider large combinations consisting of hundreds of soils concurrently and to attempt to identify and reconcile subtle differences in likelihood of occurring in each landform position for every soil in any group of soils that might possibly co-exist in space. This direct assignment of likelihood values to individual soils was considered too difficult and was not attempted. Use of intermediate likelihood values based on the 6 selected soil attributes had the added advantage of creating and preserving some record of why each soil was considered likely (or unlikely) to occur in each of the 4 landform positions. The systematic method of allocating soil attributes to landform positions based on assumed likelihood also provided a mechanism for adjusting and fine tuning the final likelihood values for individual soils without having to reconsider every soil and reassign its likelihood values in relation to all other soils.

Another alternative to the present approach might have been to base the likelihood values assigned to each soil on an analysis of available evidence on the known distribution of named soils by landform position. Such an approach would require the existence of a comprehensive electronic database containing observations and classifications of soils at known point locations. The location of the points would have to be recorded accurately using DGPS or accurate surveys. The point locations would have to be matched with a high resolution DEM classified into the 4 defined landform locations to determine the distribution of named and mapped soils by landform class. If such databases existed and if they offered sufficiently comprehensive coverage, then objective values for probability of occurrence of each named soil in each landform position could be computed and used to replace the subjective values assigned by the present procedure based on expert opinion. The currently described procedure provides a functional alternative in the absence of such comprehensive, spatially registered, databases.

The program for creating a data base linking all soils in an AGRASID map polygon to their most likely landscape position(s) (SLMSTEP2)

Hems_slm

Soilpoly	Sl_sca	Sl_lmodel	Mu_name	Lf_pos	Lf_pct	Facet_ord	New_symbol	Likelihood	Extent	Slp_50	Slp_80	Slp_len
42803501	9	U1h	PED2/U1h	UPS	25	1	LPN	77.60	20	2.0	4.0	50
42803501	9	U1h	PED2/U1h	UPS	25	1	PED	74.10	5	2.0	4.0	50
42803501	9	U1h	PED2/U1h	MID	45	2	PED	81.20	45	3.0	4.0	90
42803501	9	U1h	PED2/U1h	LOW	20	3	PED	52.90	10	2.0	3.0	40
42803501	9	U1h	PED2/U1h	LOW	20	3	ZGW	55.00	10	2.0	3.0	40
42803501	9	U1h	PED2/U1h	DEP	10	4	ZGW	92.20	10	1.0	1.0	20
42803502	9	U1h	PED1/U1h	UPS	25	1	PED	74.10	25	2.0	4.0	50
42803502	9	U1h	PED1/U1h	MID	45	2	PED	81.20	45	3.0	4.0	90
42803502	9	U1h	PED1/U1h	LOW	20	3	PED	52.90	20	2.0	3.0	40
42803502	9	U1h	PED1/U1h	DEP	10	4	PED	19.40	10	1.0	1.0	20
42803503	9	H1l	PED1/H1l	UPS	30	1	LPN	77.60	20	3.0	6.0	45
42803503	9	H1l	PED1/H1l	UPS	30	1	PED	74.10	10	3.0	6.0	45
42803503	9	H1l	PED1/H1l	MID	40	2	PED	81.20	40	4.0	6.0	60
42803503	9	H1l	PED1/H1l	LOW	20	3	PED	52.90	20	3.0	4.0	30
42803503	9	H1l	PED1/H1l	DEP	10	4	PED	19.40	10	1.0	1.0	15
42803504	9	H1m	TWS5/H1m	UPS	30	1	TWS	78.50	30	6.0	9.0	50
42803504	9	H1m	TWS5/H1m	MID	35	2	TWS	83.80	35	6.0	9.0	50
42803504	9	H1m	TWS5/H1m	LOW	25	3	TWS	48.50	5	5.0	7.0	35
42803504	9	H1m	TWS5/H1m	LOW	25	3	PED	52.90	20	5.0	7.0	35
42803504	9	H1m	TWS5/H1m	DEP	10	4	TWS	22.10	10	1.0	1.0	15

Figure 6 Illustration of the structure and content of the expanded Soil Landform Model database

The program SLMSTEP2 was run on a large subset of polygons extracted from the AGRASID digital database (xxxx polygons). It produced a new Soil Landform Model database (SLM_SLM) containing an expanded set of records for each of the original polygons in the source database. All unique soils (or non-soils) listed in the original AGRASID source file for a given polygon were linked explicitly to one or more or the 4 defined landform position classes in this new database. Each landform position was associated with the named soil or soils considered to have the highest likelihood of occurring in that landform position. Only soils named and listed in the AGRASID source file were allocated to landform positions for a given polygon. The total spatial extent of each soil within the mapped polygon was extracted from the codes in the AGRASID database for the relative extent of each named soil. Soils were allocated to landform positions until the total reported extent of each soil had been allocated to one or more of its most likely landform positions.

The procedure was almost always judged successful in allocating named soils to their most likely landform positions. The resulting allocations were deemed to be highly similar to what would have resulted had experts been asked to assign each of the listed soils to its most likely landform position or positions manually. Manual allocation of soils to landform positions for the more than 50,000 polygons in the AGRASID database would have been both prohibitively expensive and fraught with potential inconsistency. Automated allocation of soils to landform positions imposed a consistent set of rules and results. The reasons why each soil was allocated to each landform position in which it occurred could be deduced from consideration of the attributes of the soil and the values for likelihood of occurring in each landform position assigned to each attribute in the appropriate expert rule base. If soils were observed to be allocated to inappropriate landform positions, the rule bases could be examined to ascertain why and modified to achieve a more acceptable result. Each modification strengthened and improved the rule base and the understanding it contained regarding how and why different soil attributes were related to different landform positions.

The allocation procedures were both conditional and contextual. The landform position with which a soil was associated was very often conditional upon what other soils were identified in the AGRASID database as occurring in a given polygon. Thus, if eroded, thin or stony soils were identified for a given polygon, these soils were most likely to be allocated to an upper landform position and soils with orthic profiles would be shifted into mid or lower slope positions. If no eroded or thin soils were listed, soils with typical, well drained, orthic profiles were most likely to be assigned to upper landform positions. Thus, the location in the landscape to which any listed soil was assigned was very much a function of what other soils were listed as occurring in association with it. The context used to assess which landform position a given soil should be assigned to was based on consideration of both the distribution of landform classes within the current landform type and on the extent and kind of other soils associated with a soil of interest. Changing either the type of landform or the assemblage of soils under consideration affected the landform position or positions to which a given soil was most likely to be allocated.

In addition to linking each soil to its most likely landform position(s), the expanded SLM database associated a mean and limiting slope gradient and a mean slope length with each combination of soil and landform position. This association was done to facilitate potential applications involving physically based modelling or logical analysis in which it is desirable to associate a specific slope gradient or slope length with a specific soil (e.g. erosion or runoff modeling). Each landform position of each type landform was considered to possess a characteristic or defining slope gradient and slope length. Any soil allocated to a given landform position was also assumed to be under the influence of the gradient or length characteristic of the landform position.

The program SLMSTEP2 was applied to the subset of the AGRASID database a number of times and the rule bases adjusted iteratively until we were satisfied that the majority of soil landform models represented acceptable allocations of soils to landform position. It was then applied to the entire AGRASID database to compute a new, and expanded soil landform model database (SLM_SLM) for the entire region covered by the AGRASID digital soils database. It is expected that the expanded database will provide the AGRASID database with added functionality.

DISCUSSION

Did the procedures work?

The utility and efficacy of the procedures can be evaluated in two ways. In the specific case of creation of comprehensive soil-landscape models for the AGRASID database, the significant question is "did the soil-landscape models produced through application of the procedures properly reflect local expert knowledge and opinion regarding the most likely arrangement of soils by landform position for the majority of polygons in the AGRASID database?" In a general sense, one might ask whether the procedures represent a successful and practical means of capturing, codifying and applying tacit knowledge acquired and retained by local experts in soil survey.

With regard to the specific AGRASID application, the procedures should be judged in terms of how well they captured and reflected local knowledge of soil-landform relationships, not in terms of whether that tacit knowledge was consistently correct in its estimation of where soils actually occurred in the landscape. The proper test then, was not to determine through field observations whether the soils listed for any given AGRASID polygon or group of polygons actually occurred in the positions in the landscape to which they were assigned by the procedures. Rather, the proper test was to have soil survey experts review the final soil-landform models generated by application of the procedures and to assess whether the resulting allocation of soils to landform position was consistent with expectations based on their local tacit knowledge of soil-landform relationships. It was not possible to review final soil-landform models for the more than 50,000 polygons in the AGRASID database, but a significant number were checked (xxxx). Of these, the arrangement of soils by landform position was consistent with the expectations of the expert soil surveyors in almost all cases.

As indicated earlier, traditional procedures for reviewing each AGRASID polygon and manually assigning each listed soil to one or more positions in the landscape would have been prohibitively expensive to attempt for the more than 50,000 polygons. Manual procedures would also have inevitably been characterized by inconsistencies arising from differences in concepts held by different experts assigned the task of assigning soils to landform positions in any given area. Even if a single person was assigned the unenviable task of assigning soils to landform position for all 50,000 polygons, it would still have been difficult to maintain consistency for such a large number of possible combinations of soils. Tacit knowledge is not generally expressed formally and semantically (Hudson, 1992). As a result the tacit knowledge base itself can contain numerous inconsistencies and errors which will affect how an expert may place soils in the landscape but which might not be easily detected as errors. The capability of the expert rule-based procedures described here to automatically allocate all soils in the 50,000 AGRASID polygons to their most likely positions in the landscape was considered to be impressive and useful, especially considering the alternatives of manual assignment or no assignment.

In a general sense, the procedures were found to provide a very simple, elegant and effective method of capturing and applying the tacit knowledge of local soil survey experts. Hudson (1992) recognized a need for mechanisms that would permit local tacit knowledge of soil surveyors to be captured and expressed formally and semantically. In many ways the procedures described here fulfil that need. Experts are provided with a formal, quantitative and testable mechanism for capturing their understanding of where different soils are most likely to occur in the landscape and why. The process of filling in the rule base tables requires experts to express their beliefs regarding the most likely locations in which different soil characteristics or classifications will be found in a manner that is both formal and quantitative. Where soil experts can easily and readily agree on values to assign to likelihood of occurring in each of the defined landform positions for a given set of attributes or classifications, it is evident that the local knowledge base is both clearly defined and widely accepted. Cases where different experts hold widely different opinions on what values to assign to a given set of attributes or classifications indicate an inconsistency or problem in the local knowledge base and a corresponding need to revise and improve it.

Running the allocation programs provides an opportunity to apply and test the understanding contained in a rule base at any given time. If the placement of soils by landform position is inconsistent with the expectations of local experts, then a continuing problem with the rule base is indicated. The rule base can be revisited and the values corresponding to soil surveyor beliefs can be changed to determine if a more suitable outcome can be achieved if the beliefs are modified. This process provides an iterative mechanism for testing and improving the beliefs which constitute the knowledge base as formally captured by the belief tables (the rule bases).

Why did the procedures work?

The procedures were judged to be highly successful in assigning almost all soils to their most likely locations in the landscape for a large number of both simple and complex combinations of soils and landscape in the AGRASID database. This relatively high rate of success raises the question of why the procedures were so consistently able to reflect local tacit knowledge and to correctly associate soils with their most likely locations in the landscape.

The success of the procedures must be attributed to the fact that they are based on consideration of a core set of basic soil attributes and classifications whose location in the landscape is believed to be related to and controlled by physical processes influenced by landform shape and landform position. This gets back to the fundamental assumption underlying most soil surveys, which is that the location of soils in the landscape is predictable. Most soils strongly reflect the long term consequences of physical, chemical and biological processes acting on them. The major controlling variables for most of these processes are the long term status and patterns of variation in soil moisture and energy in the landscape. The movement of water over, through and under the soil is the major integrating factor in these processes. Where water flows, how it flows and where it accumulates exercise significant influence over a number of soil forming processes including erosion and deposition of soil materials, downward leaching and translocation of materials in the soil profile, upward migration of salts and soluble substances, and oxidation or reduction in the subsoil. Similarly variations in water and energy affect vegetative growth and biological activity in the soil which in turn influence rooting depth, topsoil depth, organic mater incorporation, buildup or loss and overall chemistry of the profile.

Much of the tacit expert knowledge possessed by soil surveyors reflects their understanding of how, at a local scale, topography influences the distribution and redistribution of water, energy and matter in the landscape and how this, in turn, influences soils and soil properties. The procedures for allocating soils to landform position work because they require knowledgeable soil surveyors to express and quantify this understanding. By electing to analyze this fundamental set of soil attributes in relation to landform position, the procedures were able to effectively tap into, capture and quantify the main components of local tacit knowledge possessed by a group of experienced soil surveyors.

It is unlikely that the analysis would have been as successful had we had chosen to assign a likelihood of occurring in each landform position to each soil directly, without recourse to the intermediate analysis of soil attributes. For one thing, it would have been very difficult to deal consistently and uniformly with the complexities presented by the very large number of combinations of different soils that would have had to be considered simultaneously. It is likely that manual assignment of likelihood values to each individual soil would have involved similar considerations of how topography influenced the movement and accumulation of water, materials and energy in the landscape. While similar, this consideration could never have been as formal, systematic or reproducible as that adopted for the intermediate procedures. By electing to deal with fundamental attributes of the soils (e.g. drainage, salinity, etc) we ensured that the resulting rule bases captured expert judgement on not only where in the landscape each soil was most likely to occur, but also why it was judged likely, or unlikely, to occur in any of the 4 defined landform positions. One has only to determine the characteristics of the soil and to examine the appropriate rule base to identify why a given soil has a high (or low) likelihood of occurring in a given landform position.

In essence, the allocation problem was solved by decomposing the larger problem into smaller, and more fundamental components and addressing them first individually and then collectively.

So what?

With respect to the specific application of the procedures to the AGRASID database, it is expected that the revised and expanded soil-landform model database (SLM) will prove both necessary and useful for a variety of applications. The ability to associate a specific soil in an AGRASID polygon with a specific landform position is a basic requirement for successful application of many models and decision rules.

One current proposal is to apply a physically based erosion model to the AGRASID database to assess erosion potential for all polygons in the White Area. This requires an ability to associate every soil in each polygon with a defined location in the landscape which has a defined slope gradient and slope length. Current efforts to use the AGRASID database to estimate the total amount of carbon stored in soils within the White Area and to evaluate the potential of agricultural soils in Alberta as sinks for carbon sequestration will also benefit from an ability to associate soils in AGRASID polygons with landform positions. Projects concerned with applying and evaluating precision farming technologies in Alberta have investigated using the allocation procedures to produce first approximations of detailed soil maps for section and quarter section sites for which detailed digital elevation models (DEMs) were available but for which no detailed data on soils were available. The process involved applying a recently developed model to automatically segment the DEM for a site of interest into the 4 basic landform elements (MacMillan et al., 1999). Once classified into the 4 landform elements, the soils occurring within the quarter or section of interest were identified by consulting the AGRASID database. This list of soils was then allocated to the 4 defined landform positions using the automatic allocation procedures described here. The resulting map of soils by landform element was offered as a reasonable first approximation of a high resolution soil-landform map for the site of interest.

These are just some examples of current and planned applications of the AGRASID database that benefit from an ability to associate soils with landform positions. Others are certain to exist now or to be encountered in the future. The successful capture and application of soil survey tacit knowledge to associate soils with landform position(s) can be considered to have resulted in a value-added extension to the capabilities of the original AGRASID database.

It is believed that the procedures will prove equally applicable throughout the western prairie provinces. The specific rule bases may require some minor adjustment to capture and reflect local expert knowledge in Saskatchewan and Manitoba, but the procedures themselves should be transferable. Since the process requires only data contained in the NSDB Soil Names File (SNF), there are no data restrictions, within Canada, on where it might be attempted. It may prove more difficult to apply in locales where relationships between soils and landform position are less pronounced than in the western Great Plains, but this remains to be assessed.

In the most general sense, the procedures offer a simple and effective mechanism for capturing some of the tacit knowledge retained by expert soil surveyors that has previously not proven amenable to formal expression either semantically or quantitatively. At a minimum, the procedures offer a formal, systematic protocol that facilitates efforts by experienced soil surveyors to express their understanding of where soils are most likely to occur in the landscape in a manner that is both quantitative and testable. Given the recent rapid losses of experienced soil survey personnel in most parts of Canada, it might prove worthwhile to adopt the procedures as a mechanism for recording and storing some of that vulnerable tacit knowledge and collective wisdom.

Another possible application of the procedures might be to reverse the rule bases in an attempt to predict the most likely soil or soils at any given location. Rather than starting with a known list of soils and attempting to allocate them to their most likely locations in the landscape, one might start with a known landform position and attempt to predict the most likely soil. This would require spatial data sets portraying each of the 6 basic input considerations (e.g. drainage, salinity, parent material, etc.) and might also require additional layers such as likelihood of occurrence of Solonetzic or saline soils. If these were available, or could be simulated, it is entirely conceivable that the rules might do a good job of predicting the likelihood of occurrence of any given soil in any given location.

Is that all there is?

The stark simplicity of the procedures and their almost uncanny success in capturing and applying tacit expert knowledge on soil-landform relationships raises the question of why it took so long to attempt such an exercise and why it has not been done before. Further to this point, if a significant component of the local tacit knowledge of expert soil surveyors can be captured and quantified in such a small and simple rule base, one is tempted to ask "is that all there is?".

Does this rule base truly capture a significant proportion of the tacit knowledge of soil surveyors regarding where in the landscape soils are most (or least) likely to occur. If it does, and if that is all that is really known of significance about the distribution of soils by landform position, why have soil surveyors not admitted as much and provided this limited, but admittedly useful, information to their clients and users long ago. Has it really been necessary to spend lifetimes making soil maps and acquiring "local tacit knowledge" if the largest proportion of this tacit knowledge comes down to understanding that "water flows downhill" and that soils reflect this phenomenon. Would it not have been better, as suggested by (Ellis??) to have spent the time and effort to develop and validate rules bases on the relationships of soils to landform position in a systematic and conscientious manner.

It is tempting to believe that the bulk of the "science" surrounding soil survey may have been simply "bulk", that being jargon and impenetrable codes that provided an illusion of science to a process that required only a minor amount of understanding of how a few key hydrological and pedological processes operated in landscapes. It may have been that the process of acquiring tacit knowledge through making soil maps was so enjoyable and personally rewarding that there was no incentive, or interest, in demistifying the process by explaining how elegantly simple it really was. Whatever the reason, it is the conclusion of this author that efforts put into trying to capture and express soil survey tacit knowledge formally and explicitly are long overdue and may just point the way towards mechanisms for improving mapping in the future based on systematic application of formal rule bases.

CONCLUSIONS AND RECOMMENDATIONS

- The described procedures provide a formal, systematic and reproducible mechanism for capturing, quantifying and applying the local tacit knowledge of expert soil surveyors.
- On the whole, using the final rule bases created for this project, the procedures allocated almost all of the soils in the more than 50,000 AGRASID polygons to their most likely location in the landscape, as assessed by expert opinion.
- The revised and extended Soil Landform Model (SLM) database provides more information and is of greater potential use than the original AGRASID SL and MAS files.
- The Soil Landform Model (SLM) database should be distributed as a value-added extension to the original AGRASID database (SL & MAS files).
- The NSDB Soil Names File (SNF) might benefit from having the new fields containing values for likelihood of occurring in each of the 4 landform positions added to the basic file as currently distributed.
- The provinces of Saskatchewan and Manitoba, in particular, might wish to try applying the procedures to their provincial Soil Names Files to see if they are equally applicable in these environments.
- Consideration should be given to implementing automated procedures for reformatting the data in the Soil Landform Model (SLM) file to produce effective tables and figures to assist in illustrating the patterns of distribution of soils by landform position for each AGRASID Soil-Landscape Model.

REFERENCES

Arnold, R. W. 1979. Map units test models, Cornell University. Agronomy Mimeo 79-24. Department of Agronomy, Cornell Univ. Ithaca, NY. 14853.

Arnold, R. W. 1988. Soil survey: an example of applied research, Soil Survey Horizons. 102-106.

CAESA Soil Inventory Project Work Plan Committee. 1993. Canada-Alberta Environmentally Sustainable Agriculture Agreement (CAESA) Soil Inventory Project. Revised Workplan for 1993-1994. (R. A. MacMillan, senior author). submitted to the Soil Inventory Project Technical Steering Committee. Environmental Research and Engineering Department, Alberta Research Council. Edmonton, Alberta, Canada. 22 pp.

Expert Committee on Soil Survey (ECSS), 1987. Soil Survey Handbook. Volume 1. G. M. Coen ed. Land Resource Research Centre Contribution Number 85-30. Technical Bulletin 1987 - 9E. Research Branch, Agriculture Canada.

Hudson, B. D. 1992. The soil survey as a paradigm-based science, Soil Sci. Soc. Am. J. 56: 836-841.

MacMillan, R. A. and W. W. Pettapiece. 1998. The quantitative description of landforms. Poster presentation *In:* Proceedings of the 35th Annual Alberta Soil Science Workshop. February 18 and 19, 1998, Edmonton, Alberta. pp. 254-260.

MacMillan, R. A., W. W. Pettapiece, L. D. Watson and T. W. Goddard. 1998. A landform segmentation model to support precision farming. Poster presentation *In:* Proceedings of the 35th Annual Alberta Soil Science Workshop. February 18 and 19, 1998, Edmonton, Alberta. pp. 244-253.

MacMillan, R. A., W. W. Pettapiece, L. D. Watson and T. W. Goddard. 1998. A landform segmentation model for precision farming. *In:* Proceedings of the …. Minnesota Precision Farming Conference.

MacMillan, R. A., W. W. Pettapiece, S. C. Nolan and T. W. Goddard. In press. A generic procedure for automatically segmenting landforms into landform elements using DEMs, heuristic rules and fuzzy logic. Fuzzy Sets and Systems.

MacMillan, R. A. and W. W. Pettapiece. 1997a. Soil landscape models: Automated landform characterization and generation of soil-landscape models. Technical Bulletin No. 1997-1E. Research Branch, Agriculture and Agri-Food Canada, Lethbridge, AB. 75 pp.

MacMillan, R. A. and W. W. Pettapiece. 1997b. Landform segmentation procedures manual: Step by step instructions for processing DEM data to define landform segments. Research Branch, Agriculture and Agri-Food Canada, Lethbridge, AB. 65 pp.

MacMillan, R. A. and W. W. Pettapiece. 1996. Automated generation of soil-landscape models. *In:* Proceedings of the 33rd Annual Alberta Soil Science Workshop. February 19 and 20, 1996, Edmonton, Alberta.

MacMillan, R. A. 1996a. Soil Landscape Model: Executive Summary. prepared for Agriculture and Agri-Food Canada, Alberta Land Resource Unit. 17 pp.

MacMillan, R. A. 1996b. Soil Landscape Model: Automated Morphometric Analysis. prepared for Agriculture and Agri-Food Canada, Alberta Land Resource Unit. 49 pp.

MacMillan, R. A. 1996c. Soil Landscape Model: Automated Landform Segmentation. prepared for Agriculture and Agri-Food Canada, Alberta Land Resource Unit. 14 pp.

MacMillan, R. A. 1996d. Soil Landscape Model: Automated Allocation of Soils to Landscape Elements. prepared for Agriculture and Agri-Food Canada, Alberta Land Resource Unit. 16 pp.

MacMillan, R. A. 1996e. Soil Landscape Model: Automated Production of Soil-Landscape Model Diagrams. prepared for Agriculture and Agri-Food Canada, Alberta Land Resource Unit. 11 pp.

Miller, F. P., D. E. McCormack and J. R. Talbot, 1979. Soil surveys: Review of data collection methodologies, confidence limits and utilization, ***In***: Mechanics of Track Support, Piles and Geotechnical Data. Transportation Research Board. Commission on Sociotechnical Systems. National Academy of Sciences. Washington, D.C.

A procedure for capturing and applying local tacit expert knowledge to assist in assigning soils to their most likely locations in the landscape