Land Monitor Salinity Risk Prediction
Dumbleyung and Mt Barker Regions

March 2000

A report from the NHT-funded project
Land Monitor

CMIS Task Report No. 2000/45

F. H. Evans

CSIRO Mathematical and Information Sciences

SUMMARY

This task report summarises the work and finding to date of the component of the Land Monitor project relating to salinity risk prediction.

The Land Monitor project aims to provide information about land condition, specifically salinity and the status of remnant vegetation, for the whole of the south-west agricultural region of Western Australia. It is a collaborative project involving Agriculture WA, CSIRO, CALM, DOLA, Water and Rivers Commission, the Department of Environmental Protection and Main Roads WA.

The methodology used for predicting salinity risk is as follows:

1. Use the Land Monitor DEMs and current salinity maps to create derived variables.
2. Create a decision tree that relates the derived variables to salinity risk for the ground truth areas provided.
3. Use the tree to extrapolate the decision tree model and predict risk areas.

The methodology has been tested and applied in the Dumbleyung and Mt Barker regions. The following table shows the accuracies for areas where ground data were available.

ground data area	**accuracy**	
	salt	**not salt**
Toolibin	91.5	87.5
Broomehill	70.2	92.9
Cranbrook	81.6	73.6
Tambellup	77.0	82.7
Kent	68.1	83.6
South Stirlings	85.5	39.9

The following results should be noted:

- Better results are achieved in the Dumbleyung region when only the Broomehill ground data are used.
- The best results for the Cranbrook and Tambellup ground data areas are achieved when ground data from Cranbrook, Tambellup and the Kent catchment area are combined.
- The best results for the Kent catchment area are achieved when only local ground data (ie. in the Kent) are used.
- Results in the South Stirlings area are poor. The use of regional training data produces large errors of commission but the use of local training data produces large errors of omission.

Index

1 Introduction 3
2 Salinity prediction methodology 4
3 Ground data 5
4 Feature selection 6
4.1 Feature selection procedures 6
4.1.1 Information theory approach 6
4.1.2 Forward selection 7
4.1.3 Simple forward selection 9
4.2 Results in other study areas 10
4.2.1 Forward selection 10
4.2.2 Simple forward selection 10
4.3 Regional results 11
5 Dumbleyung region risk prediction 11
5.1 Experimental results 12
6 Mt Barker Region 13
6.1 Experimental results 13
6.2 Results in the South Stirlings 15
7 Conclusions and prediction map 16
8 References 17

1 Introduction

The Land Monitor project aims to provide information about land condition, specifically salinity and the status of remnant vegetation, for the whole of the south-west agricultural region of Western Australia. It is a collaborative project involving Agriculture WA, CSIRO, CALM, DOLA, Water and Rivers Commission, the Department of Environmental Protection and Main Roads WA.

Satellite images, digital elevation maps and ground data will be used to produce maps of land condition, such as the current extent of salinity, the changes since 1990 and predictions of future salinity risk. The methods used for predicting future salinity risk were developed as part of a project funded by the Land and Water Resources Research and Development Corporation as part of the National Dryland Salinity Project (see Evans *et al.,* 1994).

The project, 'Integrating Remotely Sensed Data With Other Spatial Data Sets to Predict Areas at Risk from Salinity', investigated methods for predicting areas at risk from salinity. It focussed on methods that aimed to emulate expert-knowledge in the form of rule-based systems. A workshop held in Bunbury on June 20 - 21 in 1994 provided a forum to quantify current knowledge in explicit rule-based form. The following factors were identified as indicators of salinity risk:

- time since clearing
- depth to ground water and rate of ground water rise

- distance to existing salinity and degree of waterlogging
- climate
- geology and depth to basement
- soil type and available salt storage
- landform type: including factors like convergence, relative height, drainage density, drainage slope, valley size, position in the flow path
- historical and existing vegetation cover
- hydrogeology: vertical structures, including shears, faults, dykes, and bedrock highs and horizontal structures, including regolith stratigraphy
- land management

The project obtained all available data relating to these factors. Landsat TM satellite data were used to produce maps showing areas currently affected by salinity and areas supporting remnant vegetation. Digital elevation models (DEMs) were used to produce landform descriptors. Digital maps showing interpreted landform systems, soil systems, geology and hydrogeological structures were provided by Agriculture WA. The datasets were systematically evaluated for their ability to predict salinity risk and the final subset of input data were reduced to Landsat TM and DEM-derived variables.

The project showed that using a decision tree classifier, salinity risk areas could be predicted with 78% accuracy whilst non-risk areas were predicted with 80.5% accuracy. The predicted risk map was evaluated in the field by AgWA hydrologists with assistance from local landholders and the on-ground evaluation supported the credibility of the map.

This report describes the application of the Land Monitor method for salinity risk prediction in the Dumbleyung and Mt Barker regions, including the Salt Scenarios 2020 study area.

It is important to note a significant change in the terminology used to describe the Land Monitor prediction maps. The term *predicted risk areas* is assumed to mean areas that will have high water-tables in the future (ie. discharge areas), parts of which are likely to become saline. Whilst waterlogging may be common, predicted risk areas will not necessarily be salt-affected throughout.

2 Salinity prediction methodology

The steps used to produce the salinity risk maps are listed below. This is the standard methodology being used for Land Monitor salinity prediction.

4. The Land Monitor DEMs and current salinity maps to create the following derived variables:
    - Average upslope height
    - Average upslope slope
    - Flow slope (or steepest downhill slope)
    - Height above nearest salt
    - Height above nearest stream
    - Water accumulation (or upslope area) - and smoothed versions
    - Flow path length
    - Percentage upslope cleared area

- Total upslope cleared area

The methods used to derive the above variables are described in Caccetta (1999).

5. Create a decision tree that relates the derived variables to salinity risk for the ground truth areas provided (Evans et al., 1996).
6. Use the tree to extrapolate the decision tree model and predict risk areas.

3 Ground data

Ground data has been provided by Agriculture WA hydrologists. The ground data were digitised from interpretations of areas at risk from salinity based upon the following datasets (where available):

- rate and extent of change of salinity
- existing data on groundwater levels and the rate at which groundwater is rising
- water balance calculations (eg. AgET) to estimate the difference between current recharge, discharge from saline
- areas and base flow
- interpreted hydrogeologic data (shears, faults, dykes, bedrock highs)
- geology
- amount and type of remnant vegetation in the catchment and clearing history
- electromagnetic, radiometric and magnetic data
- ground water modelling where appropriate
- landscape processes (eg. areas with poor drainage and convergent inflow, concave inflection points (breaks of slope), non-saline discharge areas)

The interpretations were then digitised and converted into raster format with values of 1 where there is no risk from salinity and 2 where future salinity risk occurs.

Ground data were provided for the following areas: Broomehill, Toolibin, Cranbrook, Tambellup, Kent and the South Stirlings.

Of these areas, Broomehill and Toolibin are located in the Dumbleyung region and the remainder in the Mt Barker region. Figure 1 shows the ground truth map for the Toolibin catchment.

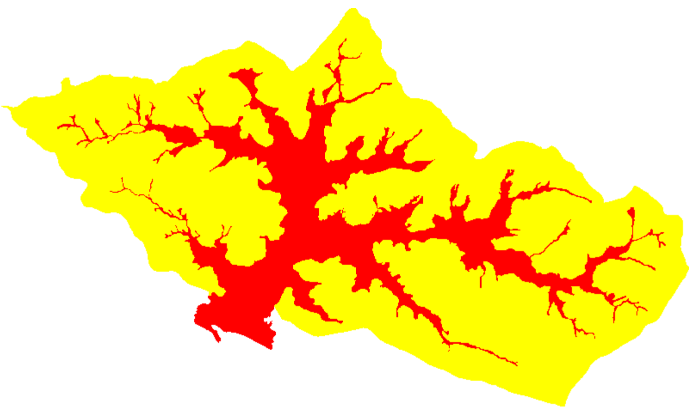

Figure 1 Ground truth map for Toolibin. Risk areas are shown in red.

4 Feature selection

This section describes the results of using feature selection procedures for determining the optimal subset of DEM-derived variables for predicting salinity risk areas. Preliminary studies were conducted using data for the Broomehill study area and assessed in the Toolibin, Kent, Cranbrook and Tambellup areas.

The results described in this report have been produced using 50% of the available ground data for training and the remaining 50% for testing. Options -c0 and -m5 are used to train the decision trees. The results for individual study areas have been produced by training and testing only on local data (ie. data from that study area). Section 4.3 presents the regional results obtained when the training and test data are combined for the broader Dumbleyung and Mt Barker regions.

4.1 Feature selection procedures

We use feature selection to determine whether salinity risk can be accurately mapped using a subset of the nine DEM-derived variables listed above. By reducing the number of variables, we aim to reduce the size of the decision tree (rendering it simpler to interpret) and improve the accuracy and visual outlook of the risk maps that are produced.

Exhaustive feature selection involves testing each possible subset of the variables and assessing the subsets according to their accuracy on a set of independent test data. The primary disadvantage of the exhaustive approach is the large number of subsets that must be tested (in this case, 9!). Instead, we have used three heuristic approaches: the information theory approach, the forward selection approach and the simple forward selection approach.

4.1.1 Information theory approach

The c5.0 induction algorithm uses information theory to evaluate splits. Two splitting criteria are implemented. The variable that the decision tree splits on is determined by calculating the expected information generated that is useful for classification by each possible split on each variable. Thus, the variable on which the decision tree makes its first split is the variable that provides the most useful information for salinity risk prediction.

The information theory approach is as follows:

(i) Produce a decision tree using all variables.

(ii) Examine the decision tree to determine the variable on which it first splits.

(iii) Remove this variable.

(iv) Repeat until no variables remain.

Table 1 shows the results of using this approach.

Table 1 Information theory approach

number of variables	first split	salt accuracy	not salt accuracy
9	UpClr	70.1	92.9

8	wa	71.1	92.1
7	FpLen	68.6	92.6
6	Habove Streams	68.5	92.8
5	Avg Height	68.1	91.5
4	Habove Salt	62.6	92.9
3	FlSlope	56.3	93.1
2	PerClr	51.5	88.3
1	AvSlp	4.6	97.4

Table 1 shows the variables that provide the most information for salinity risk prediction as assessed using the information theory approach. Forward selection is applied by:

(i) Produce a decision tree using the variable that provides the most information.

(ii) Add the variable that provides the next largest amount of information.

(iii) Continue until no significant gains in accuracy are obtained.

The results shown in Table 2 contain an inconsistency that provides some doubt about the efficacy of this approach - when the second most important variable is added, the salt accuracy is decreased. It is for this reason that the second heuristic approach to feature selection has been applied in section 4.1.2.

Table 2 Forward selection using the information theory approach

number of variables	variable added	salt accuracy	not salt accuracy
1	UpClr	57.9	94.4
2	wa	55.1	95.2
3	FpLen	57.5	94.5
4	Habove Streams	67.3	94.1
5	Avg Height	67.8	93.5
6	Habove Salt	69.7	93.1
9	FLSlope, PerClr, Avg Slope	70.1	92.9

4.1.2 Forward selection

The forward selection approach is as follows:

(i) Determine the best single variable.

(ii) Add the variable that improves the accuracy by the greatest amount.

(iii) Continue until no significant gains in accuracy are obtained.

Table 3 shows the results of using this approach. Variables added are shown in bold.

Table 3 Forward selection

number of variables	variable(s)	salt accuracy	not salt accuracy
1	**Habove Salt**	60.4	90.6
1	UpClr	57.9	94.4
1	Habove Streams	55.9	95.6
1	Avg Height	55.8	95.4
1	wa	55.1	95.1
1	FpLen	50.6	95.1
1	FlSlope	47.1	92.0
1	PerClr	41.4	87.3
1	Avg Slope	4.6	97.4
2	**+ FpLen**	65.7	92.3
2	+ wa	65.5	93.7
2	+ Avg Height	65.1	93.5
2	+ UpClr	65.0	94.3
2	+ PerClr	61.3	91.1
2	+ FlSlope	61.1	92.6
2	+ Avg Slope	59.4	92.0
2	+ Habove Streams	58.5	93.7
3	**+ Habove Streams**	70.2	89.6
3	+ Avg Slope	68.1	92.5
3	+ wa	67.0	92.0
3	+ Avg Height	66.8	91.6
3	+ UpClr	66.8	93.0
3	+FlSlope	66.2	92.2
3	+ PerClr	64.2	93.1
4	**+ FlSlope**	71.1	92.2
4	+ UpClr	71.0	93.1
4	+ wa	70.3	92.5

4	+ Avg Height	68.2	93.1
4	+ Avg Slope	67.9	92.9
4	+ PerClr	67.6	92.9
5	+ UpClr	70.5	93.0
5	+ PerClr	70.1	92.0
5	+ wa	69.8	92.2
5	+ Avg Slope	69.3	92.3
5	+ Avg Height	69.2	92.7

Note that:

- The addition of a fifth variables reduces the accuracy with which salt is mapped.
- The accuracy achieved using only four variables is greater (by 1%) than that achieved when all nine variables are used.

The optimal subset of four variables is: flow slope, height above nearest salt, height above nearest stream and flow path length.

4.1.3 Simple forward selection

The simple forward selection approach is as follows:

(i) Determine the best single variable.

(ii) Add the next best single variable.

(iii) If no significant gain in accuracy is made, eliminate that variable and add the next best single variable.

(iv) Continue until no significant gains in accuracy are obtained.

Table 4 shows the results of using this approach. Variables added are shown in bold.

Table 4 Simple forward selection

number of variables	variable(s)	salt accuracy	not salt accuracy
1	**Habove Salt**	60.4	90.6
2	**+ UpClr**	65.0	94.3
3	**+ Habove Streams**	68.5	93.2
4	+ Avg Height	68.6	92.9
4	**+ wa**	69.1	92.9
5	**+ FpLen**	69.8	92.2

The optimal subset of five variables is: height above salt, uplsope cleared area, height above stream, smoothed water accumulation and flow path length. The unsmoothed water accumulation variable should be included in areas where only subsections of the broad valley floor are at risk.

4.2 Results in other study areas

4.2.1 Forward selection

Tables 5-8 show the results achieved by the forward selection approach in the Toolibin, Cranbrook, Tambellup and Kent River study areas. An extra variable, the unsmoothed water accumulation, has been included in the analyses for study areas in the Mt Barker region. The reason for this becomes evident in section 4.2.2.

Table 5 Toolibin results

number of variables	salt accuracy	not salt accuracy
9	86.8	93.5
4	87.9	92.6

Table 6 Cranbrook results

number of variables	salt accuracy	not salt accuracy
10	55.1	81.0
4	48.7	90.0

Table 7 Tambellup results

number of variables	salt accuracy	not salt accuracy
10	63.9	88.4
4	68.9	86.0

Table 8 Kent results

number of variables	salt accuracy	not salt accuracy
10	79.2	80.2
4	77.2	77.9

Note that:

- In all study areas except Cranbrook, the use of the four variables selected using the forward selection approach provides results that are similar or improved upon those achieved when all of the variables are used.

4.2.2 Simple forward selection

Tables 9-12 show the results achieved by the simple forward selection approach in the Toolibin, Cranbrook, Tambellup and Kent River study areas.

Table 9 Toolibin results

number of variables	salt accuracy	not salt accuracy
9	86.8	93.5
5	86.5	92.7

Table 10 Cranbrook results

number of variables	salt accuracy	not salt accuracy
10	55.1	81.0
5	57.5	87.0

Table 11 Tambellup results

number of variables	salt accuracy	not salt accuracy
10	63.9	88.4
5	69.4	87.5

Table 12 Kent results

number of variables	salt accuracy	not salt accuracy
10	79.2	80.2
5	80.1	77.7

Note that:

- The use of the 5 variables selected using the simple forward selection approach provides results that are similar or improved upon those achieved when all of the variables are used.
- The only study area in which the results are poor (whilst being improved) is Cranbrook. Visual inspection of the risk maps suggested that the accuracy might be improved with the inclusion of the unsmoothed water accumulation variable. This is supported by the accuracies shown in Table 13.

Table 13 Cranbrook results

number of variables	salt accuracy	not salt accuracy
10	55.1	81.0
6	67.0	85.3

4.3 Regional results

Tables 14 and 15 show the results achieved by the simple forward selection approach in the Dumbleyung and Mt Barker regions.

Table 14 Dumbleyung results

number of variables	salt accuracy	not salt accuracy
9	79.2	94.4
5	80.7	94.2

Table 15 Mt Barker results

number of variables	salt accuracy	not salt accuracy
10	71.7	83.7
5	71.3	84.0
6	69.9	84.6

5 Dumbleyung region risk prediction

This section describes the salinity prediction for the Dumbleyung region. The geographic region is:

Top left: 495070E 6382725N

Bottom right: 642275E 6236750N

The results show that the best overall decision tree is that produced when only the Broomehill training data is used.

5.1 Experimental results

The following results are a sub-sample of experiments that were performed to ensure that the decision tree prepared would produce results that could be extrapolated to map risk areas for the Dumbleyung region. All of the following risk maps and statistics were produced using the decision tree classifier c5.0 with options -c0 -m5, and costs set such that errors for risk areas incur double the cost of errors in non-risk areas. The use of costs is required because there are significantly more non-risk training sites than risk training sites.

Table 16 shows the the preliminary results using the original nine DEM-derived data sets as supplied and the subset of variables as determined using the feature selection techniques described in section 4. In the table, global fitting used the combined Toolibin and Broomehill training data, while local fitting used only training data from the tabled catchment.

Table 16 Dumbleyung Risk Prediction - Summary of Analyses

Method	Overall		Toolibin		Broomehill	
	salt	not salt	salt	not salt	salt	not salt
global fitting 9 variables	79.2	94.4	87.7	94.5	67.0	95.5
local fitting 9 variables	-	-	86.8	93.5	70.2	92.9
Toolibin training data only	-	-	86.8	93.5	53.9	95.1
Broomehill training data only	88.4	88.6	91.5	87.5	70.2	92.9
global fitting 5 variables[1]	80.7	94.2	88.7	93.5	68.2	95.2
local fitting 5 variables	-	-	86.5	92.7	69.8	92.2

The highest accuracy for predicting risk areas in both Broomehill and Toolibin occurred when the Broomehill ground truth data were used to train the decision tree classifier. The decision tree produced using these data was retained and used to produce the prediction map (version 1). Whilst mapping risk areas with sufficient accuracy, the prediction map tends to over-estimate the risk areas (by 7.1% in Broomehill and 12.5% in Toolibin). This problem can be reduced by combining the Broomehill and Toolibin ground truth data, but the accuracy with which risk areas can be mapped is then reduced (from 70.2% to 67.0% in Broomehill and from 91.5% to 87.7% in Toolibin). Whilst the reduction in accuracy is small, visual inspection of the maps produced in this manner shows many areas within valleys and streamlines that are missed by the risk map.

The results also showed that whilst the Broomehill ground data can be extrapolated to accurately map risk areas in Toolibin, a decision tree trained using the Toolibin ground data will only map 53.9% of the risk areas in Broomehill.

[1] See simple forward feature selection, section 4.2.2

6 Mt Barker Region

This report describes the salinity prediction for the Mt Barker region. The geographic region is:

Top left:	677400E	6111000N
Bottom right:	192600E	181400N

The results show that:

1. The best overall decision tree (global fit) is that with accuracies shown in Table 3.
2. Results in the South Stirlings and higher rainfall zones suggest that these areas should also be modelled separately as risk areas are being over-estimated.

6.1 Experimental results

The following results are a sub-sample of experiments that were performed to ensure that the decision tree prepared would produce results that could be extrapolated to map risk areas for the Mt Barker region. All of the following risk maps and statistics were produced using the decision tree classifier c5.0 with options -c0 -m5, and costs set such that errors for risk areas incur double the cost of errors in non-risk areas. The use of costs is required because there are significantly more non-risk training sites than risk training sites.

Table 17 shows the the preliminary results using the original nine DEM-derived data sets. Examination of the errors that could be seen in the resulting prediction maps showed over-estimation of risk in broad valley areas. For this reason, a ten-variable set that also included the un-smoothed water accumulation variable was also tested. Smaller subsets of variables were selected based upon previous work that used feature selection techniques to reduce the number of variables.

In each of the following tables, global fitting used the combined training data from each subcatchment, while local fitting used only training data from the tabled catchment.

Table 17 Summary of Analyses

Method	Overall		Cranbrook		Tambellup		Kent	
	salt	not salt	salt	not salt	salt	not salt	salt	not salt
global fitting 10 variables	71.7	83.7	71.1	78.4	70.6	87.3	72.5	81.8
local fitting 5 variables[2]	-	-	57.5	87.0	69.4	87.5	80.1	77.7
local fitting 6 variables	-	-	67.0	85.3	68.2	88.3	81.2	79.0
global fitting 5 variables	71.3	84.0	71.8	77.9	70.9	87.1	71.5	83.1
global fitting 6 variables	69.9	84.6	67.0	85.3	66.7	88.9	81.2	79.1

Further examination of the resulting prediction maps and the input variables showed some errors in the *height above salt* variable - salinity maps for areas within the Mt Barker Landsat TM scene had been excluded from the data processing.

Table 18 shows the amended results after this error was corrected.

Table 18 Summary of Analyses

Method	Overall		Cranbrook		Tambellup		Kent	
	salt	not salt	salt	not salt	salt	not salt	salt	not salt
global fitting	73.2	79.8	78.6	68.5	77.0	82.7	68.1	83.6
local fitting	-	-	58.1	88.9	71.7	85.3	79.8	78.2

Table 19 shows the results of modifying the *height above streams* variable to include a greater proportion of local valley systems. The global decision tree provides the best overall results produced so far.

[2] See simple forward feature selection, section 4.2.2

Table 19 Summary of Analyses

Method	Overall		Cranbrook		Tambellup		Kent	
	salt	not salt	salt	not salt	salt	not salt	salt	not salt
global fitting	76.1	81.1	81.6	73.6	76.7	84.5	73.2	84.0
local fitting	-	-	67.9	82.7	78.3	82.6	77.1	79.3

It should be noted that:

1. Decision trees produced using just the Cranbrook training data perform poorly - this is a small area and there is insufficient data to accurately map salinity in Cranbrook using this data alone.
2. Decision trees produced using just the Kent training data perform better than the global decision trees. This suggests that the Kent catchment should be modelled separately.

For these reasons, Table 20 shows the results of using just the Cranbrook and Tambellup training data to predict risk in these catchments. The overall accuracies are calculated using test data from the Cranbrook and Tambellup catchments combined.

Table 20 Summary of Analyses

Method	Overall		Cranbrook		Tambellup	
	salt	not salt	salt	not salt	salt	not salt
global fitting	73.2	82.3	78.3	74.0	71.3	85.8

The results show that the accuracies for both the combined data and for Cranbrook and Tambellup data are lower than when the Kent data is included. This suggests the possibility that the Kent training data extrapolates well to the outer area. The final prediction map is thus produced using the global decision tree trained using the combined ground data from the Kent, Cranbrook and Tambellup catchments, with accuracies as shown in Table 19.

6.2 Results in the South Stirlings

At this stage, the South Stirlings data were digitised and the models were tested in the South Stirlings area. The results are shown in Table 21. The global model over-estimated salinity risk in the South Stirlings with around sixty percent of the non-risk areas being mapped as risk areas. Local models also performed poorly performed, mapping only 39% of the risk areas correctly. These results suggest that:

1. The South Stirlings should be modelled separately to the remainder if the Mt Barker scene.
2. Additional training data is required to accurately map the South Stirlings area.

Note that the accuracy with which risk areas are mapped using only local training data is not improved beyond 50% even when the costs are changed from 2:1 to 10:1.

Table 21 South Stirlings Analyses

Method	South Stirlings	
	salt	not salt
global fitting	85.5	39.9
local fitting	39.0	95.9

Examination of the extrapolated risk maps south of this point suggests that risk areas are also being over-estimated in the higher rainfall zones and therefore additional training data are required for these areas.

7 Conclusions and prediction map

The prediction map for the Dumbleyung region has been produced using the decision tree classifier trained on only the Broomehill data with accuracies presented in Table 16. The accuracies achieved in the two areas with ground data were acceptable and further ground-truthing has been performed in the Boscabel and Towerinning areas. Detailed assessments of the salinity risk maps at specific geographically located areas were noted for both the Boscabel and Woodanilling areas. In general:

1. The overall impression of the risk maps is that they are good.
2. Errors in local valleys are of two types. Some should show less risk extending up the valley and some should show more.
3. Broad valley floors are exclusively mapped as risk. However, there are many areas where the groundwater is near-surface (ie. 1-2m) within the broad valleys that may be slightly higher in elevation than the true valley floor (less than 1m height differences can be significant) that will not become salt-affected. These areas may currently be supporting a healthy crop or pasture cover.

The third point has caused a significant change in the terminology used to describe the Land Monitor prediction maps. The term *predicted risk areas* is assumed to mean areas that will have high water-tables in the future (ie. discharge areas), parts of which are likely to become saline. Whilst waterlogging may be common, predicted risk areas will not necessarily be salt-affected throughout.

The prediction map for the Mt Barker region has been produced using the decision tree classifier with accuracies presented in Table 19. The input attributes used were:

- Flow slope
- Height above salt
- Height above streams
- Smoothed water accumulation
- Unsmoothed water accumulation
- Upslope cleared area

Whilst the results are acceptable in the Kent, Cranbrook and Tambellup catchments, risk areas are severely over-estimated in the South Stirlings and southern higher rainfall areas. More training data are required to correct these problems. The current version of the map for these areas is not to be released under any circumstances.

It should also be noted that in the North Stirlings area where the accurate DEMs were not available and 10m contour data were used in their place, the majority of the broad valley floors have been mapped as risk areas. This problem may be fixed when the accurate DEMs become available.

8 References

Caccetta, P. C. (1999), *Some methods for deriving variables from digital elevation models for the purpose of analysis, partitioning of terrain and providing decision support for what-if scenarios*, CMIS internal task report, available from http://www.cmis.csiro.au/rsm/research/dems/demsfeb1999.htm

Evans, F. H. (2000), *Dumbleyung salinity risk prediction – version 1*, CMIS internal task report, Land Monitor project.

Evans, F. H. (2000), *Feature selection for salinity risk prediction*, CMIS internal task report, Land Monitor project.

Evans, F. H. (2000), *Mt Barker salinity risk prediction*, CMIS internal task report, Land Monitor project.

Evans, F. H. (2000), *Dumbleyung salinity risk prediction – version 2*, CMIS internal task report, Land Monitor project.

Evans, F. H., Caccetta, P. C., Ferdowsian, R., Kiiveri, H. T. and Campbell, N. A. (1995), *Predicting salinity in the Upper Kent River catchment*, A report from the LWRRDC project *Integrating remotely sensed data with other spatial data sets to predict areas at risk from salinity*.

Evans, F. H., Caccetta, P. C., and Ferdowsian, R. (1996), 'Integrating remotely sensed data with other spatial data sets to predict areas at risk from salinity'*, Proceedings of the 8th Australasian Remote Sensing Conference*, available on cdrom.