P. A. Whigham and P. F. Crapper

This chapter describes the application of a grammatically-based Genetic Programming system to discover rainfall-runoff relationships for two vastly different catchments. A context-free grammar is used to define the search space for the mathematical language used to express the evolving programs. A daily time series of rainfall-runoff is used to train the evolving population. A deterministic lumped parameter model, based on the unit hydrograph, is compared with the results of the evolved models on an independent data set. The favourable results of the Genetic Programming approach show that machine learning techniques are potentially a useful tool for developing hydrological models, especially when the relationship between rainfall and runoff is poor.

5.1 Introduction

Many problems of interest to natural resource scientists may be expressed in the form of a time series model. This chapter describes a geographic problem, represented at the catchment scale and daily time scale, which attempts to relate the rainfall incident on a catchment to the stream flow at the exit from the catchment. The produced function has the form $\mathcal{F}(r_1,r_2,..r_n)$, where r_1 .. r_n are rainfall variables from current and/or previous days. $\mathcal{F}$ represents the current days streamflow. The difficulties involved with producing $\mathcal{F}$ are clear when we consider the numerous complex features that directly and indirectly contribute to the measured behaviour. The response of the catchment (especially Australian catchments) is highly capricious, depending not only on the catchment characteristics (e.g. topography, area), vegetation characteristics and antecedent conditions, but the meteorological conditions (e.g. areal distribution of rainfall) in a highly non-linear and unpredictable fashion. Developing models that describe this relationship help in understanding the overall behaviour of the catchment and support the development of more process-based models and catchment classification schemes. Additionally, many natural resource models use streamflow as an input. Although reliable rainfall records often exist for a catchment, only limited streamflow data is generally available.

This chapter compares the performance of a variant of Genetic Programming (GP) (Koza, 1992) to a traditional hydrological model which predicts the rainfall-runoff relationship for specific catchments. Rainfall-runoff models have been previously developed using other machine learning techniques, such as neural networks (Minns and Hall,1996), however the form of these models was not easily translated into interpretations of catchment process and behaviour. The advantage of using a symbolic learning system, such as GP, is that the resulting model may be interpreted in terms of the process and behaviour of the catchment. This can often help in understanding the underlying processes that drive the catchment response and can be used to classify and generalise differing catchments. GP has been previously applied to time series prediction (Mulloy, Riolo and Savit, 1996) however no previous work, to the authors knowledge, has applied GP to the problem of predicting rainfall-runoff relationships.

5.2 The Genetic Programming System CFG-GP

A variant of genetic programming, context-free grammar GP (CFG-GP), (Whigham, 1996) uses a grammar to define the space of legal sentences that can be explored during evolution. The system allows a transparent definition of language bias by instantiating the way in which terminals of the language may be legally combined.

A formal grammar is a production system which defines how nonterminal symbols may be transformed to create terminal sentences of a language. A grammar is represented by a four-tuple (N, Σ, P, S), where N is the alphabet of nonterminal symbols, Σ is the alphabet of terminal symbols, P is the set of productions and S is the designated start symbol. For example, the following grammar, G_{math}, defines a language for generating all possible mathematical expressions using the operators +, -, *, /, and a set of random real numbers, represented by the symbol, $\Re$.

$$
\begin{aligned}
G_{math} = \\
&\{S, \\
&N = \{M\}, \\
&\Sigma = \{+,-,*,/,\Re\}, \\
&P = \\
&\quad \{ \\
&\quad\quad S \rightarrow M \\
&\quad\quad M \rightarrow +\,M\,M \mid -\,M\,M \mid *\,M\,M \mid /\,M\,M \mid \Re \\
&\quad \} \\
&\}
\end{aligned}
$$

The initial random population is generated using the grammar by selecting random productions which match the current nonterminals in the derivation, starting with S, and limited by some maximum depth of derivation tree. A derivation step represents the application of a production to some string which contains a nonterminal. In general, a series of derivation steps may be represented by a syntax tree or derivation tree. These trees have genetic operators applied to them in a manner similar to normal GP program trees, except that crossover sites between derivation trees must use matching nonterminal sites.

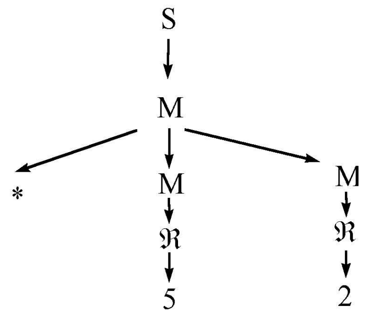

Figure 5.1
A Derivation Tree for the expression * 5 2, generated using the grammar G_{math}. S,M and $\Re$ are nonterminals symbols, while *, 5 and 2 are terminal symbols of the language.

For example, using G_{math}, the expression string * 5 2 could be generated from the following derivation steps:

$$S \Rightarrow M \Rightarrow * \, M \, M \Rightarrow * \, \Re \, M \Rightarrow * \, 5 \, M \Rightarrow * \, 5 \, \Re \Rightarrow * \, 5 \, 2 \qquad \text{Equation 5.1}$$

The corresponding derivation tree for this sequence is shown in Figure 5.1. Crossover and mutation operators are applied directly to these trees.

Two search operators are used to modify the evolving rainfall-runoff model; the crossover operator is used as the search operator for each generation, mixing elements of potentially useful partial solutions in an attempt to build a better solution; a hill-climbing mutation is used as a fine-tuning operator that allows the random constants within the final best solution to be modified in an attempt to move towards a more optimal solution.

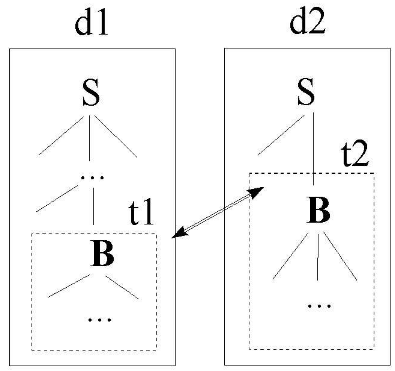

Figure 5.2
Crossover between derivation trees d1 and d2. The subtrees t1 and t2, with root node nonterminal B, are swapped to produce two new derivation trees (i.e. new programs). Crossover always swaps subtrees with the same root nonterminal.

Crossover applies 2 (parent) individuals and creates 2 (offspring) individuals. Each crossover operation is defined by two parameters: the probability of crossover occurring and the nonterminal $B \in N$ where crossover will be applied. Assuming that the crossover operator has been selected, two programs with derivation trees d_1 and d_2 are selected from the current population using a proportional fitness selection. Crossover, as shown in Figure 5.2, is then performed as follows:

1. Randomly select a subtree t1 from d1 with root node B.
2. Randomly select a subtree t2 from d2 with root node B.
3. Swap t1 and t2 thereby creating two new derivation trees $d1^*$ and $d2^*$.
4. Insert $d1^*$ and $d2^*$ into the next-generation population.

If d_1 or d_2 do not contain the nonterminal B then no crossover is possible and the operation is aborted. The benefit of using derivation trees to represent the population now becomes clear; by defining crossover to swap subtrees at the same nonterminal guarantees that the space of possible programs is constrained to be part of the language defined by the grammar.

The rainfall-runoff grammar G_{flow} (see Section 5.4) defines mathematical expressions which are initially seeded with variables relating to current or previous rainfall, and random real numbers. These real numbers are used as constants throughout the evolution, and are combined into the partial solutions that are evolved. A final solution that uses one or more of these constants may be improved (based on the training data) by slight modifications of these constants. The hill climbing mutation applies small random changes to the constants of the best final program (derivation tree) and maintains the new solution only if it improves the final performance based on the training data. This mutation is applied a fixed number of times and may be considered as a fine tuning of the evolved solution.

A proportional fitness measure is used to select programs (i.e. derivation trees) each generation for crossover and reproduction.

5.3 Rainfall Runoff Modelling

One of the traditional approaches to hydrograph modelling (Jakeman, Littlewood, and Whitehead, 1990) is to use the concept of the Instantaneous Unit Hydrograph (IUH). The IUH is defined as the hydrograph produced by the instantaneous application of a unit depth of rainfall to a catchment. The shape of the IUH is similar to a single peak hydrograph with a rapid rise and a slower decay. The fundamental assumption in the IUH model is that the precipitation input is equal to the integrated streamflow output. The non-linear relationship between rainfall and streamflow has led to the development of the concept of *effective rainfall*, which is determined by applying a non-linear filter to the raw rainfall data. This effective rainfall is then equated with the integrated streamflow for the specified catchment.

The IHACRES model applied in this paper is based on IUH principles. The model defines a unit hydrograph for total streamflow by defining separate unit hydrographs for the quickflow and the slowflow components. The model is defined by six parameters, four of which are determined directly from the raw rainfall, streamflow and temperature (or a surrogate), while the other two (the non-linear parameters) are calibrated using a trial and error search procedure, optimising the model to fit the observed rainfall-runoff relationship. The fundamental conceptualisation in the non-linear module of the model is that catchment wetness varies with recent past rainfall and with evapotranspiration. A 'catchment wetness index' is computed for each time step on the basis of recent rainfall and temperature. The percentage of rainfall which becomes effective rainfall in any time step varies linearly between 0% and 100% as the catchment wetness index varies between zero and unity. An alternative conceptualisation of the catchment wetness index is that it represents the proportion of the catchment at a given time step which contributes eventually to streamflow, but it is important not stretch the physical interpretation of catchment wetness index too far. Conceptualisation of spatially distributed processes in both the non-linear and the linear modules is severely restricted by the spatially lumped nature of the model. An advantage of this approach is that the model needs only a small number of parameters. Additional details about the model are contained in (Littlewood and Jakeman, 1994). It is worth noting that this model is considered to be one of the standard approaches to rainfall-runoff modelling, and has been used successfully for a number of years.

5.4 CFG-GP Setup

The grammar, G_{flow}, used by CFG-GP to develop the rainfall-runoff models allowed simple mathematical functions to be evolved, and was defined as follows.

G_{flow} =
{S,
N = {EQU, NL, EXPN},
Σ={+,-,*,/,exp,r0,r1,r2,r3,r4,r5,av5,av10,av15,
av20,av25, av30, av40, av50,av60,av100,$\Re$},
P =
{
$S \rightarrow$ + EQU NL
$NL \rightarrow$ * EQU EXPN
$EXPN \rightarrow$ exp EQU
$EQU \rightarrow$ + EQU EQU | - EQU EQU
$EQU \rightarrow$ * EQU EQU | / EQU EQU
$EQU \rightarrow$ exp EQU
$EQU \rightarrow$ r0 | r1 | r2 | r3 | r4 | r5
$EQU \rightarrow$ av5 | av10 | av15 | av20 | av25
$EQU \rightarrow$ av30 | av40 |av50 | av60 | av100
$EQU \rightarrow \Re$
}
}

The terminal symbols r0,r1,..r5 represent the rainfall for the current day up to the last 5 days rain. The av5, av10, .. av100 terminals are the average rainfall for the last 5, 10, ... 100 days, respectively. The terminal $\Re$ is a random floating point number between -10.0 and 10.0 which is generated for each occurrence of $\Re$ when the initial population is created. These random constants are potentially modified using the hill climbing mutation once the best, final generation, solution has been found using crossover. The exponential function is represented by the terminal string "exp". The grammar has a structural bias to form equations that are composed of a linear component and a non-linear (exponential) component. This is shown by the production S $\rightarrow$ + EQU NL, which forces all programs to have the minimal structure of A + B * exp(C), where A, B and C are climate variables or a random real number. The production EQU $\rightarrow$ exp EQU allows the exponential function to be included with any part of the evolved mathematical expression. The language bias merely forces the use of the exponential function at least once in the final solution.

Table 5.1
CFG-GP Parameter Settings used for all experiments.

CFG-GP Parameter	Value
POPULATION SIZE	1000
GENERATIONS	50
GRAMMAR	G_{flow}
MAX. TREE DEPTH	15
CROSSOVER $\otimes$ = {EQU}	90%
HILL CLIMB MUTATION	1000 times
FITNESS MEASURE	Minimise RMSE

Table 5.1 shows the CFG-GP setup parameters which were used to develop both catchment models. The crossover operator is applied only to the nonterminal EQU, with a probability of 90%. Hence approximately 10% of the population is passed unchanged into each subsequent generation. This ensures that good solutions are not prematurely removed from the population and that the building blocks that are useful are maintained. The root mean square error (RMSE) was used as the fitness measure. If P_t is the predicted runoff value at time t, A_t is the actual runoff value at time t, and there are N points in the training data (N > 1), then RMSE is defined as follows (Chatfield, 1984):

$$RMSE = \sqrt{\frac{\sum (P_t - A_t)^2}{N - 1}} \qquad \text{Equation 5.2}$$

The CFG-GP system used the same training data as IHACRES to evolve the rainfall-runoff models. For IHACRES, the training data was used to calibrate the constants which appear in the IHACRES conceptual model. In the case of CFG-GP, the training data was used to both evolve suitable constants and to develop the underlying structure of the model itself. Comparison of results will only refer to the simulation (test) runs which use previously unseen data for the same location. The CFG-GP system was run 100 times for each catchment, with the best equation on the training data selected as a candidate for the final solution. The best candidate solution on the unseen data was selected as the resulting equation. The RMSE is used to compare the IHACRES and CFG-GP models for the simulation period. As an additional measure of performance the error in predicted total discharge (i.e. the sum of streamflow) for the simulation period is also calculated. However, this error is not used as part of the CFG-GP fitness function.

5.5 Catchment Descriptions and Results

In order to test the modelling approaches two very different catchments were chosen. The first catchment was the Teifi catchment at Glan Teifi in Wales, United Kingdom. The second catchment was located within the Namoi River catchment in northeastern New South Wales, Australia.

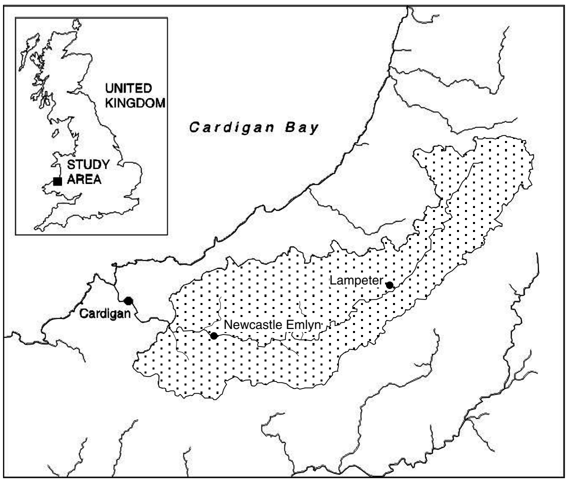

Figure 5.3
Location of Glan Teifi Catchment.

5.5.1 The Glan Teifi Catchment

The Teifi catchment is a rural catchment draining 893.6 km^2 with an average annual rainfall of 1368 mm. This station was maintained and operated by the UK Environment Agency and the data can be obtained from the Institute of Hydrology, Wallingford, Oxon, U.K. Compared with the Namoi catchment, the number of rain days per annum is much greater at Teifi but the maximum daily rainfall is only about half the value for the Namoi. The other major difference between the catchments is that runoff percentages (i.e. total runoff/total rainfall * 100) are very much higher at Glan Teifi. . The calibration run for Teifi was done from 27th July, 1982 to 31st July, 1985 and the simulation run was done from 23rd July, 1979 to the 27th July, 1982. For the calibration period the runoff percentage was 66.7% and for the simulation period the runoff percentage was 74.95%. The measured rainfall and streamflow for the Teifi catchment between July 1979 and July 1982 is shown in Figure 5.4 (rainfall events are shown as black columns). It is worth noting that the Glan Teifi catchment has a strong seasonal signal and that there appears to be a strong relationship between rainfall and runoff.

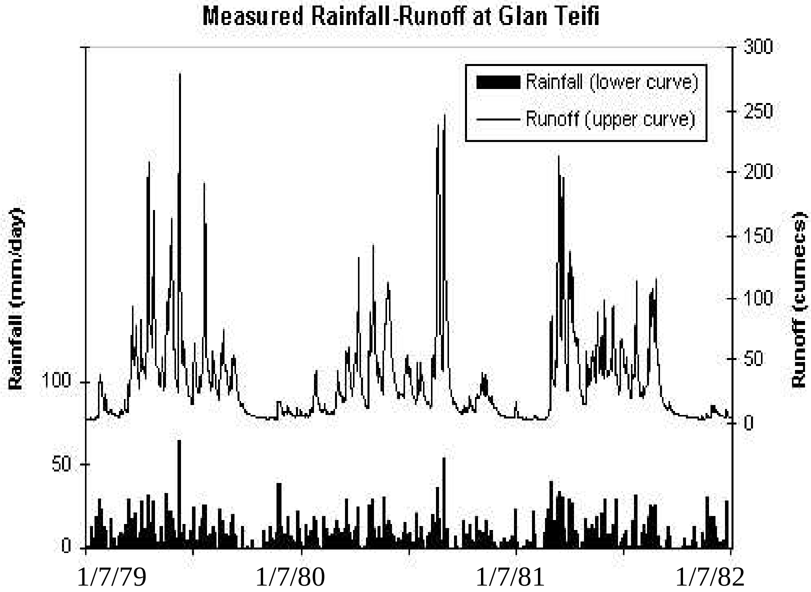

Figure 5.4
Measured Rainfall and Runoff at Glan Teifi.

5.5.2 Results

The simulated streamflows determined by IHACRES and CFG-GP are shown in Figures 5.5 and 5.6. The daily error for each model, calculated as (predicted flow - actual flow), is shown in the right-hand graphs for each figure. A visual comparison with the measured streamflow (Figure 5.4), indicates that both approaches have captured the basic response of the catchment, however IHACRES appears to have better represented the extreme streamflow events. The root mean square error (RMSE) for IHACRES was 0.0139 and for CFG-GP was 0.0142. The total discharge for the simulation period was measured as 35,600 cumecs, with IHACRES predicting 34,776 cumecs (2.3% error) and the evolved CFG-GP solution predicting 33,295 cumecs (6.4% error). The CFG-GP equation which was evolved for the Teifi catchment was defined as follows.

$$+(+(r1,+(av40,*(av10,av100))),*(*(av5, +(-17.121983,*(av5,av40))),exp(-3.739896))) \qquad \text{Equation 5.3}$$

This may be simplified to give:

$$r1 + av40 + (av10 * av100) + (av5 * (-17.121983 + (av5 * av40)))* 0.0237 \qquad \text{Equation 5.4}$$

Equation 5.4 shows that the catchment was influenced by antecedent conditions that could extend for several months into the past (the av100 variable represents the average rainfall for the last 100 days). Note also that the constant exponential expression in Equation 5.3, namely exp(-3.739896), means that the resultant equation for runoff is only a linear function of r1 (previous days rain) ,av5, av10, av40, and av100. There is no non-linear component.

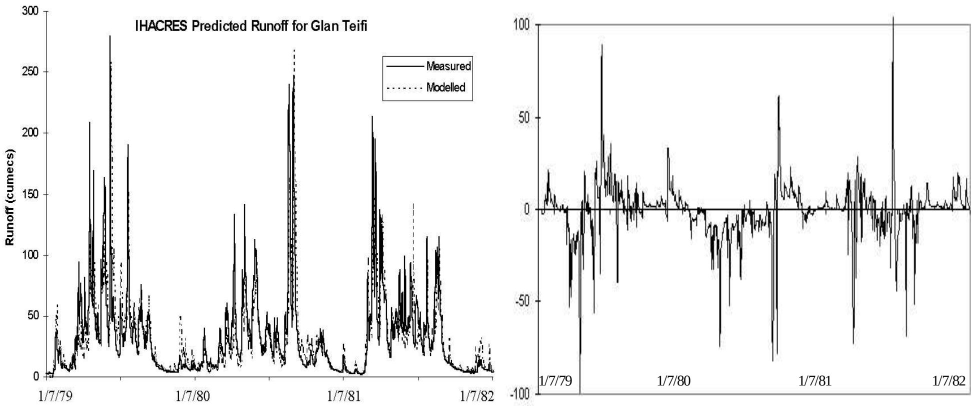

Figure 5.5
IHACRES Modelled Runoff at Glan Teifi for the test (unseen) time period. The Error (Predicted Value - Measured Value) is shown in the right-hand graph.

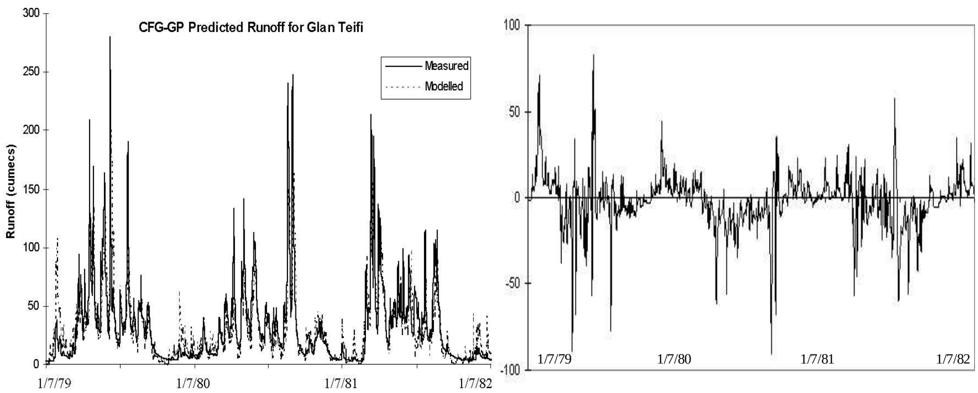

Figure 5.6
CFG-GP Modelled Runoff at Glan Teifi for the test (unseen) time period. The Error (Predicted Value - Measured Value) is shown in the right-hand graph.

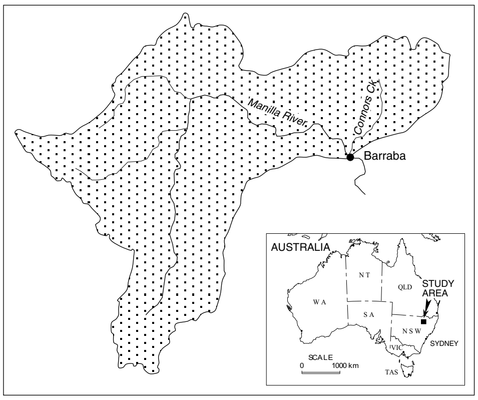

Figure 5.7
The Manilla River at Barraba, N.S.W.

5.5.3 The Namoi River Catchment

The Namoi River catchment (see Figure 5.7) was chosen to be as different as possible from the Teifi catchment. Using the Department of Land and Water Conservation (the gauging authority) naming convention the catchment is referred to as 419030 or the Manilla Rv at Barraba (30° 23’ 24” S and 150° 37’ 08” E). This catchment is approx. 568 km^2 and drains the southern part of the Nandewar Range. Within the catchment there are three reliable long-term raingauge stations with average annual rainfalls of 686mm, 704mm and 727mm. These stations have a reasonable spread of location and altitude and the average of the three values has been used as the catchment rainfall. More sophisticated techniques do exist for determining catchment rainfall but given that there were only three rainstations such sophisticated approaches are inappropriate. For very large rainfall events (greater than 100mm) there was a strong relationship between rainfall and runoff but as the size of the event decreased the relationship between rainfall and runoff became more random. The rainfall in this part of the country is strongly summer dominated, which influenced the selection of the calibration and simulation periods. The calibration run was done from 13 November 1965 to 10 March 1966 and the simulation run was done from 4 November 1966 to 13 March 1967. This short time period for calibration was necessary because IHACRES could not converge when longer periods were chosen. This was due to the requirement that the start and end points should be selected to be at low flow periods. When the calibration period was chosen with several low flow periods this assumption was violated and meant that the model would not converge. In spite of this catchment having a comparatively high rainfall (by Australian standards at least) and our selection of the high rainfall months, the runoff percentage over the calibration period was only 6.12% and the simulation period was 8.24%.

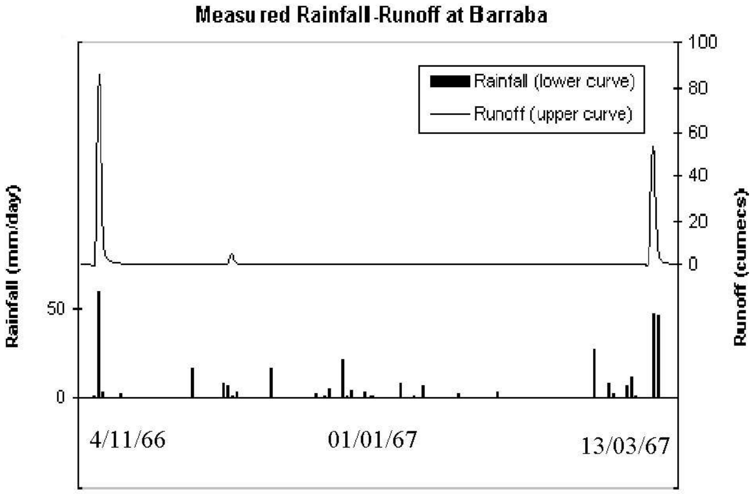

Figure 5.8
Measured Rainfall and Runoff at Barraba.

5.5.4 Results

The measured rainfall and subsequent streamflow for the simulation period in the Barrada catchment is shown in Figure 5.8. As can be seen from this data, for large events there is a strong relationship between rainfall and runoff. For smaller events, however, there is not a significant relationship between the rainfall and runoff.
The simulated streamflows determined by the two approaches are shown in Figures 5.9 and 5.10. The error for each model (predicted value - measured value) is shown in the graphs on the right-hand side of each figure. The RMSE for the IHACRES approach was 0.0474 and for CFG-GP was 0.0439. The total discharge for the simulation period was measured as 187 cumecs, with IHACRES predicting 330 cumecs (76% error) and the evolved CFG-GP solution predicting 99 cumecs (47% error). For the purposes of our comparison (based on RMSE) these results are similar.

The evolved equation found by CFG-GP was:

+(/(/(/(/(r0,-1.911790),-0.474622),-4.164400),
-(/(-1.542888,5.944119),-(av10,r0))),
*(r0,exp(/(0.251564,av10)))) Equation 5.5

which may be simplified to

$$(r0/-3.7) - 0.26 - av10 - r0 - (r0 * \exp(0.251564 / av10)) \quad \text{Equation 5.6}$$

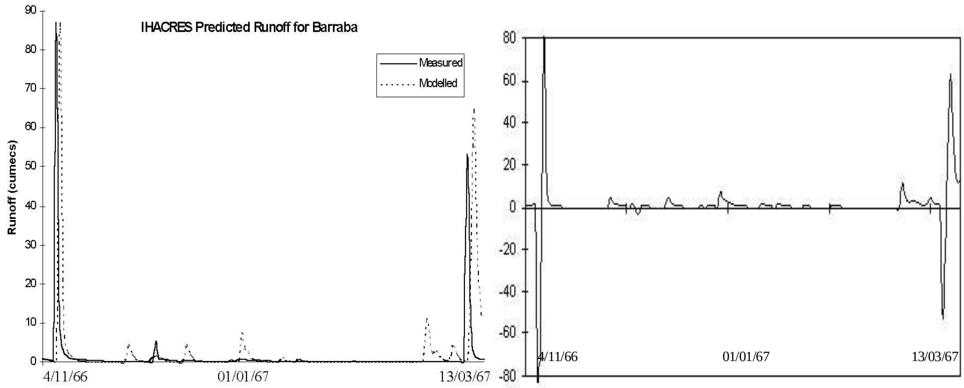

Figure 5.9
IHACRES Modelled Runoff at Barraba for the test (unseen) time period. The Error (Predicted - Measured Value) is shown in the right-hand graph.

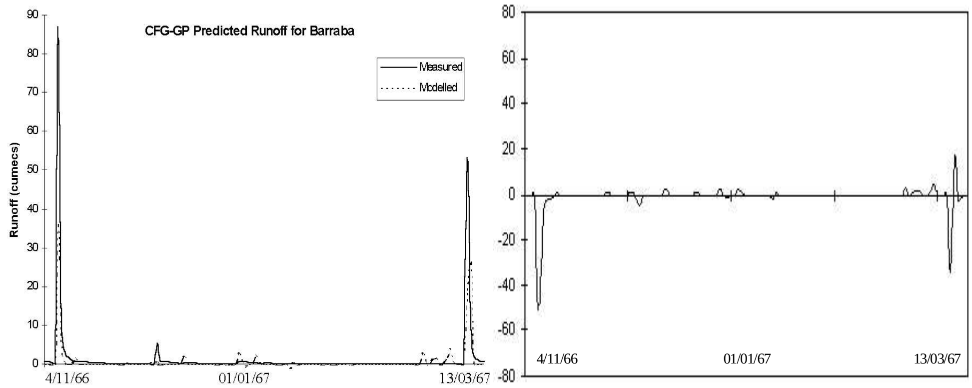

Figure 5.10
CFG-GP Modelled Runoff at Barraba for the test (unseen) time period. The Error (Predicted - Measured Value) is shown in the right-hand graph.

It is worth noting that Equation 5.6 uses the current days rainfall (r0), and the average of the last 10 days rainfall (av10). Additionally, Equation 5.6 has the nonlinear term (exp(/(0.251564,av10)), which is a function of av10. A comparison of equations 5.4 and 5.6 shows that the two catchments have been modelled in very different ways. The Welsh catchment has been modelled using long term averages in a linear combination, whereas the Australian catchment has been modelled using short average times and the current day in a nonlinear fashion. This would suggest that the underlying processes that are driving the water movement throughout both catchments are quite different.

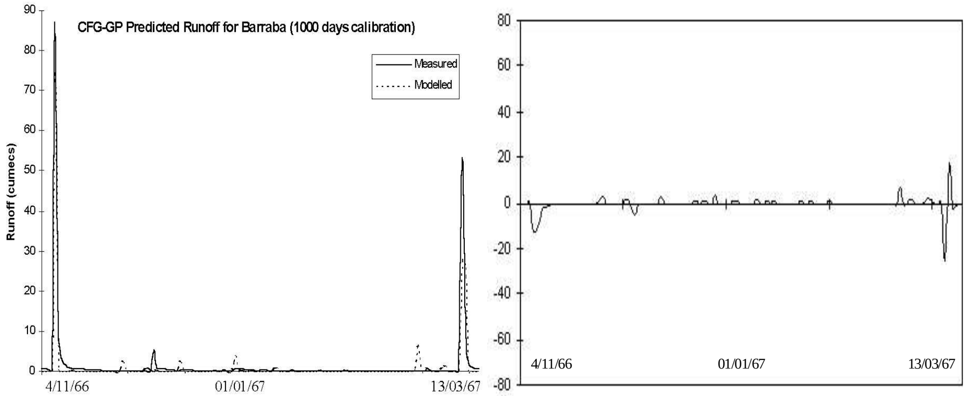

Figure 5.11
CFG-GP Modelled Runoff after using 1000 days of training data, applied to the previous 3 months test (unseen) time period. The Error (Predicted - Measured Value) is shown in the right-hand graph.

When an attempt was made to calibrate over different consecutive seasons for the Barraba data, the IHACRES model was not able to find coefficients to suit all seasons, and therefore could not converge. This accounts for the short calibration and simulation periods of only 4 months that has been used for testing these models. However, the CFG-GP approach, because it makes no assumption about underlying relationships, was able to be calibrated over successive seasons and therefore use more information about the catchment response to rainfall. When CFG-GP was calibrated using a period of 1000 days the resultant model achieved significantly better results on the original simulation data set (RMSE = 0.0237). Additionally, the predicted total discharge changed to 157 cumecs (16% error), which is superior to either previous solution. The response of this modelled streamflow, and the associated errors, are shown in Figure 5.11. The evolved equation was:

+(exp(+(/(exp(-4.874963),-0.796608),
/(/(r0,-1.864706),+(-3.018028,-4.418388)
))),*(-3.240420,exp(-1.181253))) Equation 5.7

which may be simplified to give

$$\exp(0.0096 + (r0/13.35)) - 0.994 \qquad \text{Equation 5.8}$$

The interesting comparison between Equations 5.6 and 5.8 is that using the larger dataset for calibration (training) resulted in a solution that was a nonlinear function solely of (r0), which represents the current days measured rainfall. No average rainfall value was found to be useful. This implies that the Barrada catchment has a very quick response between rainfall and runoff, and no significant seasonal signal.

5.7 Discussion

The previous examples of applying CFG-GP to modelling rainfall-runoff has been encouraging. The use of a simple, non-linear mathematical grammar has allowed the system to produce equations that capture some measure of the underlying response of the catchment. The linear, strongly seasonal model evolved for the Teifi catchment has a natural interpretation with the underlying climatic and topographic characteristics of this catchment. The non-linear, weakly seasonal model evolved for the Namoi catchment also corresponds with the perceived behaviour of this Australian landscape.

The grammar, G_{flow}, had only a weak bias towards forming certain types of mathematical expressions. Future work will involve extending the set of useful mathematical functions (power and logarithmic functions are often used in natural system modelling) and exploring other language forms which may have more direct interpretation with natural processes.

5.8 Conclusion

In the present work we have compared the results obtained with a deterministic lumped parameter model, based on the unit hydrograph approach, with those obtained using a stochastic machine learning model.

For the Welsh catchment the results between the two models were similar. Since rainfall and runoff were highly correlated the deterministic assumption underlying the IHACRES model was satisfied. Therefore IHACRES could achieve a satisfactory correlation between calibration and simulation data. It is also interesting to note that for this catchment the runoff ratio was approximately 70% which suggests that a relationship does indeed exist between the rainfall and runoff. The CFG-GP approach does not require any causal relationships but achieved similar results.

The behaviour of the studied Australian catchment was found to be quite different from the Welsh catchment. The runoff ratio was very low (7%) and hence the *a priori* assumptions of IHACRES (and other deterministic models) were a poor representation of the real world. This was demonstrated by the inability of IHACRES to use more than one seasons data for calibration purposes and only able to use data from a high rainfall period. Since the CFG-GP approach did not make any assumptions about the underlying physical processes, calibration periods over more than one season could be used. These led to significantly improved generalisations for the modelled behaviour of the catchment.

In summary, either approach worked satisfactorily when rainfall and runoff were correlated. However, when this correlation was poor, the CFG-GP had some advantages because it did not assume any underlying relationships. In these circumstances the use of evolutionary algorithms warrants further consideration.

Acknowledgments

The authors would like to thank Heinz Buettikofer of CSIRO Land and Water for producing the artwork for each of the maps. The authors would also like to thank the Institute of Hydrology, Wallingford, Oxon, UK for supplying the spatial and rainfall/runoff data of the Glan Teifi catchment.

References

Chatfield, C. (1984). The Analysis of Time Series: An Introduction. Chapman and Hall, 1984

Jakeman, A. J., Littlewood, I. G. and Whitehead, P.G. (1990). Computation of the instantaneous unit hydrograph and identifiable component flows with application to two small upland catchments. Journal of Hydrology, 117, 275-300, 1990

Koza, J.R. (1992). *Genetic Programming: on the programming of computers by means of natural selection*. A Bradford Book, MIT Press, 1992

Littlewood, I. G. and Jakeman, A. J. (1994) A new method of rainfall-runoff modelling and its applications in catchment hydrology. In: P. Zannetti (ed) Environmental Modelling (Volume II), Computational Mechanics Publications, Southampton, UK, 143-171.

Minns, A W, and Hall, M J, (1996) Artificial neural networks as rainfall-runoff models. Hydrological Sciences Journal,41(3), pp 399-417

Mulloy, B., Riolo, R. and Savit, R. (1996). Dynamics of Genetic Programming and Chaotic Time Series Prediction. In Genetic Programming: proceedings of the first annual conference. J. Koza (ed). Stanford University, CA, USA: MIT Press, July 1996

Whigham, P.A. (1996) Language Bias, Search Bias and Genetic Programming. In Genetic Programming: proceedings of the first annual conference. J. Koza (ed). Stanford University, CA, USA: MIT Press, July 1996