A hybrid neural networks and numerical models approach for predicting groundwater abstraction impacts

S. J. Birkinshaw, G. Parkin and Z. Rao

ABSTRACT

A rapid assessment method for evaluating the impacts of groundwater abstraction on river flow depletion has been developed and tested. A hybrid approach was taken, in which a neural network model was used to mimic the results from numerical simulations of interactions between groundwater and rivers using the SHETRAN integrated catchment modelling system. The use of a numerical model ensures self-consistent relationships between input and output data which have a physical basis and are smooth and free of noise. The model simulations required large number of input parameters and several types of time series and spatial output data representing river flow depletions and groundwater drawdown. An orthogonal array technique was used to select parameter values from the multi-dimensional parameter space, providing an efficient design for the neural network training as the datasets are reasonably independent. The efficiency of the neural network model was also improved by a data reduction approach involving fitting curves to the outputs from the numerical model without significant loss of information. It was found that the use of these techniques were essential to develop a feasible method of providing rapid access to the results of detailed process-based simulations using neural networks.

Key words | artificial neural networks, groundwater, groundwater abstraction impacts, numerical modelling, river–aquifer interactions

S. J. Birkinshaw (corresponding author)
G. Parkin
School of Civil Engineering and Geosciences,
University of Newcastle upon Tyne,
Newcastle NE1 7RU,
UK
Tel.: +44 191 2226 319
E-mail: *s.j.birkinshaw@ncl.ac.uk*

Z. Rao
Halcrow Group Ltd,
Burderop Park,
Swindon SN4 0QD,
UK

INTRODUCTION

In order to assess the impact of groundwater abstractions on river flows the normal approach is to use either analytical models or numerical models. Analytical solutions (Theis 1941; Hantush 1959; Hunt 1999) are limited in their accuracy or applicability due to the models' assumptions of homogeneous, isotropic aquifer systems of infinite extent. Many analytical models can represent gaining or losing rivers which are in hydraulic connection with aquifers, but it is difficult to obtain solutions for such nonlinear problems as disconnected rivers, transient (seasonal) recharge inputs or interactions with floodplain wetlands. The usual approach to increase the realism in models is to use a numerical modelling approach. Most numerical models of river–aquifer interaction (see Parkin *et al.* (2007) for recent examples) involve solution of equations for surface water

routing and groundwater flow, with coupling between the two models usually based on a simple Darcy calculation (Winter 1995). However, numerical modelling is both time-consuming and expensive. Hence, an "intermediate" technique is desirable. The approach developed in this study is to use artificial neural networks (ANNs) to mimic numerical model simulations of generic river–aquifer systems, providing a hybrid system which retains the complexity of numerical models and the speed of analytical models.

Artificial neural networks are now widely used in hydrology. They have been used in rainfall–runoff modeling (Tokar & Johnson 1999; Anctil & Lauzon 2004; Kumar *et al.* 2005) and groundwater hydrology (Balkhair 2002; Shigidi & Garcia 2003; Daliakopoulos *et al.* 2005). The use of hybrid models combining numerical models and ANNs

is less common. Recent examples, including Dibike & Abbott (1999), Ochoa-Rivera *et al.* (2002) and Chua & Holz (2005), have focused mostly on river flows, although hybrid models have also been used to mimic process-based models for groundwater remediation and river water quality (Rao & Jamieson 1997; Rao & O'Connell 1999).

This paper describes the application of a hybrid numerical and ANN model to evaluate the impact of groundwater abstractions on river flows. As far as the authors know this is the first time this approach has been used for this type of problem. The study was carried out for the Environment Agency of England and Wales, and was designed to be applicable to as wide a range of river–aquifer and abstraction scenarios as possible. The procedure involves carrying out a large set of numerical model simulations for several generic hypothetical case studies representing most river–aquifer settings identified in England and Wales, with outputs for time series of flow depletions, depletions along a river channel and water table drawdowns around the abstraction well. The ANN is then trained using input–output data from the numerical model, which allows the rapid simulation of the effect of groundwater abstractions on river flows. The full hydrogeological background to the study, details of the numerical model simulations, tests against field datasets and the full parameter sets for all settings are described in Parkin *et al.* (2007). This paper focuses on the methods of data processing and neural network training used in the study, which are demonstrated through their application to one of the settings. Only the relevant datasets and parameters used for this setting are considered here.

METHOD

In the first stage of the study, river–aquifer systems were classified into three generic settings: (1) a major chalk or sandstone aquifer in direct contact with the river; (2) a major chalk or sandstone aquifer with a shallow gravel aquifer in direct contact with the river; and (3) a shallow gravel aquifer in direct contact with the river. This work focuses only on the shallow valley aquifer as a case study to demonstrate the methodology. Details of the other settings are given in Parkin *et al.* (2007). The shallow aquifer hydrogeological setting represents those instances in which

a sand and gravel aquifer is the only aquifer in communication with a given reach of a river. Such is the case in the lower reaches of the Middle Thames Valley, for instance, where the Middle Thames Gravels are underlain by London Clay, yet nevertheless support major public supply abstractions in their own right. The methodology adopted in this study for assessing the impact of groundwater abstraction on river flow depletion for the shallow sand and gravel aquifer can be summarized as follows:

1. *Determine parameters and values*: the "hydrogeological setting" is used as the basis for defining the input parameters and their ranges, together with the required outputs.
2. *Define simulations*: the numerical model simulations are designed based on a generic model for this setting, using input parameter values chosen from a range of values.
3. *Run models*: the numerical model is run for each set of input data, producing a set of output data
4. *Train ANN model*: the ANN is trained ("calibrated") using the input–output data, and tested ("validated") against an independent set of numerical model results, to demonstrate that it is capable of reproducing the simulation results.
5. *Use trained ANN model for predictions*: once trained, the ANN can be used for predictions, within the range of its training data.

Figure 1 shows the procedure for running the simulations and training the ANN model as a flow chart. A schematic of the ANN can also be seen in Figure 2. This shows how three types of input data are considered (aquifer parameters, abstraction information and recharge details) and three types of output data (time series of flow depletion, spatial pattern of flow depletion along the river channel and spatial pattern of aquifer drawdowns).

A wide range of plausible values for the valley sand and gravel aquifer underlain by low permeability strata can be seen in Table 1, which shows nine different input parameters. This includes three aquifer parameters (gravel aquifer hydraulic conductivity, specific yield and bed sediment hydraulic conductivity), three parameters defining the abstraction (distance of the borehole from the river, abstraction rate and abstraction duration) and three parameters defining the recharge (mean annual recharge,
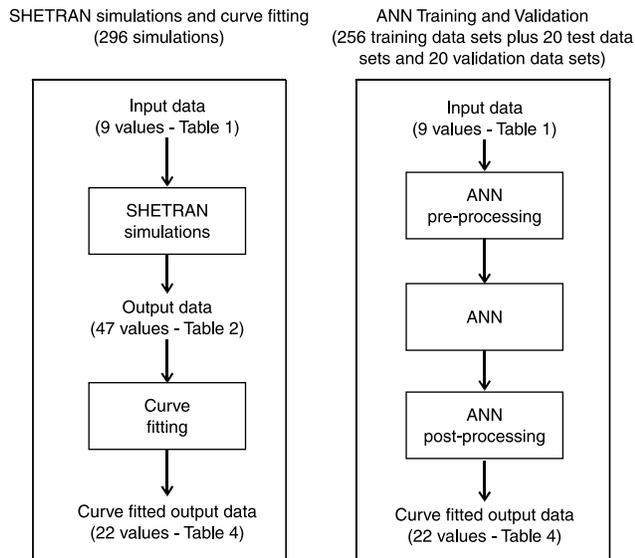
SHETRAN simulations and curve fitting
(296 simulations)

Input data
(9 values - Table 1)

↓

SHETRAN
simulations

↓

Output data
(47 values - Table 2)

↓

Curve
fitting

↓

Curve fitted output data
(22 values - Table 4)

ANN Training and Validation
(256 training data sets plus 20 test data
sets and 20 validation data sets)

Input data
(9 values - Table 1)

↓

ANN
pre-processing

↓

ANN

↓

ANN
post-processing

↓

Curve fitted output data
(22 values - Table 4)

Figure 1 | Flow chart showing the procedure for running SHETRAN and training the ANN for the shallow gravel river–aquifer setting.

recharge seasonality and time from the start of the abstraction to the maximum recharge). For each parameter four values were selected: for example, for the abstraction duration the values chosen were 1 d, 20 d, 60 d and 365 d. To run a full set of numerical model simulations with all combinations of each of these parameter values would have resulted in 262,144 simulations, an unrealistically large number. A subset of these possible simulations was defined using an orthogonal array approach (Hedayat *et al.* 1999), a structured combinatorial method which ensures that as much of the "input parameter space" as possible was covered with the simulations. In this study the abstraction rate was considered separately, but for the other eight parameters the orthogonal array was OA(64,8,4,2). This means that there were 64 combinations of 8 parameters with each parameter taking one of four values (Table 1)
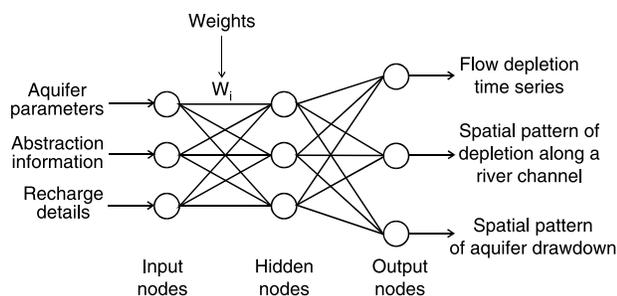
Weights

Aquifer
parameters

W_i

Abstraction
information

Recharge
details

Flow depletion
time series

Spatial pattern of
depletion along a
river channel

Spatial pattern
of aquifer drawdown

Input
nodes

Hidden
nodes

Output
nodes

Figure 2 | Schematic of the ANN.

with a strength of 2. The strength of 2 means that, for any two of the parameter values, all combinations of the parameters were used an equal number of times. For each of the 64 combinations a simulation was run for each of four abstraction rates. In addition 20 test simulation and 20 validation simulations were carried out. In total, this approach reduced the total number of simulations in this case study to 296.

In general, any model could be used which is capable of representing the processes which are considered to be important. In this study, the physically based spatially distributed SHETRAN model was used (Ewen *et al.* 2000), because of its capability of representing integrated surface and subsurface flows. Finite difference methods are used to solve the partial differential equations for flow and transport that are at the heart of SHETRAN. As such, the catchment area is discretized into rectangular computational elements, and the underlying soil zone and aquifer are represented by columns of cells extending downwards from each of the surface grid squares. The river network is represented by a network of links which run around the edge of the ground surface link elements. A SHETRAN model was set up for a generic region covering a 2 km reach of a river running through the centre of a 1 km wide valley. A typical SHETRAN grid resolution was used, varying in size from 5 m near the stream channel to 100 m at the edge of the valley. A flow boundary condition was defined at the upstream end of the river, and no-flow boundary conditions were used for all aquifer boundaries. Figure 3 shows the SHETRAN set-up. Input to the SHETRAN model then consists of information about recharge and groundwater abstractions together with parameter values for the type of aquifer (Table 1). The simulations were transient with time-varying recharge. Important aquifer properties are the hydraulic conductivity, which is a measure of how fast water can flow through the soil, and the specific yield, which is the fraction of a saturated block of soil that can be abstracted. Output from the model (Table 2) consists of three types of data: firstly, time series of flow depletions at the catchment outlet; secondly, spatial patterns of flow depletions along the river channel and thirdly, water table drawdown at various points around the abstraction well (Figure 3). The spatial output is produced at two separate times: at the end of the abstraction and at the time of

Table 1 | SHETRAN input data (t_s is the start of the abstraction and t_r is the time of maximum recharge)

Symbol	Description	Units	Values
D	Distance of borehole from river	m	25, 112, 262, 500
Q	Abstraction rate	m³/d	500, 1000, 2000, 5000
t_{ss}	Time from to t_s to t_r (days from 0–365)	d	0, 90, 180, 270
t_d	Duration of abstraction	d	1, 20, 60, 365
K_v	Gravel aquifer hydraulic conductivity	m/d	1, 4, 50, 600
S_y	Gravel aquifer specific yield	–	0.1, 0.2, 0.3, 0.5
K_b	River bed sediment hyd. conductivity	m/d	0.00004, 0.04, 4, 400
R	Mean annual recharge	mm/yr	0, 200, 400, 1000
R_s	Recharge seasonality	–	0, 0.3, 0.6, 1

maximum flow depletion. The output values were all calculated by comparing the simulation results with those from a corresponding simulation in which there is zero groundwater abstraction.

SHETRAN MODELLING

Using the orthogonal array approach discussed previously, 296 simulations were set up to run over a 25 yr period, for a

Figure 3 | SHETRAN generic model for this case study. (a) Plan view, (b) Cross section.

systematically determined subset of all possible parameter values. The example in Table 3 gives the values of the parameters for a validation set, for which the SHETRAN depletions can be seen in Figure 4. Time series for flow depletions from SHETRAN at the catchment outlet (q_t in Table 2) were given at a set of discrete times (1, 2, 5, 10, 20, 50, 100 and 200 d after the start of the abstraction and at 0, 1, 2, 5, 10, 20, 50, 100 and 200 d and 1, 2, 3, 4, 5, 10 and 25 yr from the end of the abstraction). The values focused on the time periods at the start and end of the abstraction as these periods generally have the most rapid change in depletion rates.

Spatial output was produced for flow depletions along the river channel and drawdown depth around the well at the end of the abstraction and at the time of maximum depletion (Figure 3). For depletions along the river channel values were output from SHETRAN at the following points compared to the distance (D) between the borehole and the stream: downstream D, downstream $\frac{1}{2}D$, at the nearest point of the river to the borehole, upstream $\frac{1}{2}D$ and upstream D (Figure 3). At a distance $4D$ downstream from the nearest point of the river to the borehole, the flow depletion from these simulations is at its maximum and therefore represents the total flow depletion (i.e. the value in the flow depletion time series). The pattern of drawdowns around the well output from SHETRAN were: at a distance of $D/2$ from the abstraction borehole away from the river, parallel to the river and towards the river; at the nearest point of the river to the borehole; and at a distance of $D/2$ from the river on the opposite side of the river to the borehole (Figure 3).
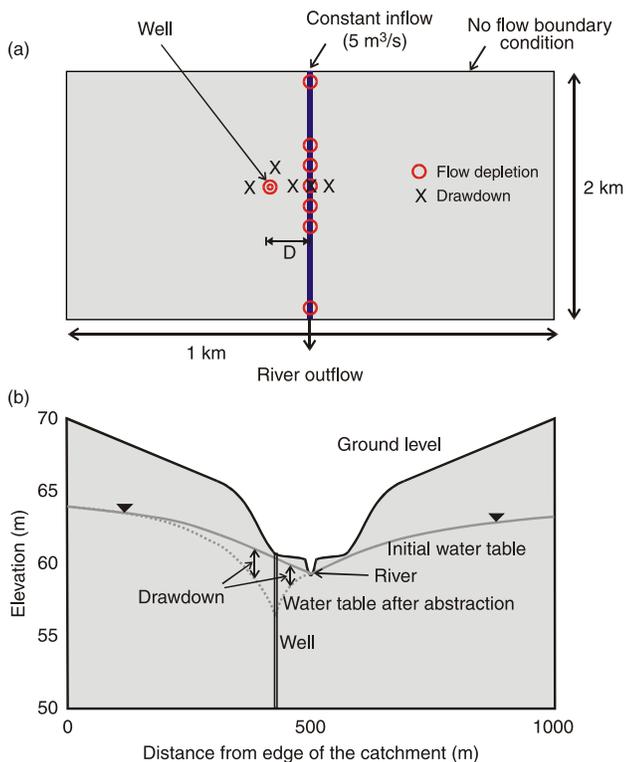
Table 2 | Output data from SHETRAN

Symbol	Description	Units	No. of values	Variable number
q_t	Flow depletion time series	m^3/d	24	1–24
q_x	Depletion profile in river	m^3/d	10	25–34
d	Aquifer drawdown	m	10	35–44
t_{max}	Time of maximum flow depletion	d	1	45
q_{max}	Maximum flow depletion	m^3/d	1	46
d_w	Drawdown in the well	m	1	47

DATA REDUCTION

Direct training of the ANN using SHETRAN input data
and output data was not appropriate for two reasons.
Firstly, this would have required too many weights for
the number of training sets. Secondly, the flow depletion
time series and the spatial pattern of flow depletion are
composed of a series of points whereas in fact they are part
of a continuous response from SHETRAN, and the points
are correlated. This section discusses how the number of
degrees of freedom in the input–output relationships was
reduced by parametrizing some aspects of the outputs with
negligible loss of information content. This represents a
form of data reduction which can improve the effective-
ness of a neural network model (Dawson & Wilby 2001).

The outputs from the SHETRAN model for the flow
depletion time series and for flow depletions along a river
channel form a family of well-behaved curves. For example,
every time series of the flow depletion starts at zero at the
start of the abstraction and then climbs smoothly and

monotonically to its maximum before falling monotonically
back towards zero. This demonstrates an advantage of the
hybrid approach, as the output from numerical models is
noise-free, avoiding the need to filter the data (Dawson &
Wilby 2001). A generalized functional form based on an
S-shaped hypsometric curve was chosen to fit to the
SHETRAN output, where a and p are the shape parameters:

$$f(x) = \left[1 - (1 - x)^{\frac{1}{a}} \right]^{(p-1)} \tag{1}$$

This is a suitable curve as $f(0) = 0$ and $f(1) = 1$, but between
these points the curve climbs smoothly and monotonically
with a wide variety of shapes possible. A total of four curves
were fitted to each of the outputs from the SHETRAN
simulations (Figure 5). Curves 1 and 2 represent the time
series of total flow depletion in the river, the first up to the
time of maximum depletion and the second from the time of
maximum depletion to the end of the simulation. Curves 3
and 4 represent the spatial extent of flow depletion along
the river at two specific times (at the end of the abstraction
and at the time of maximum depletion, t_{max} in Table 2).

The data to fit the flow depletion curve are obtained from
values 1–24 (q), 45 (t_{max}) and 46 (q_{max}) in the SHETRAN
output variables (Table 2). For curve 1, the SHETRAN output
consists of paired data of the form: $(t_{1,i}\ q_{1,i})$, $i = 1, n_1$, where
$t_{1,i}$ and $q_{1,i}$ are the time and the corresponding flow depletions
for the n_1 data points up to the time of maximum depletion.
The form of the hypsometric equation for the first fitted curve
is then

$$\frac{q_{1,j}}{q_{max}} = \left[1 - \left(1 - \frac{t_{1,j}}{t_{max}} \right)^{\frac{1}{a_1}} \right]^{(p_1-1)} \tag{2}$$

where a_1 and p_1 are the curve fitting parameters, selected
to produce the best fit to this equation using the

Table 3 | Validation input data used in the example (t_s is the start of the abstraction and t_r is the time of maximum recharge)

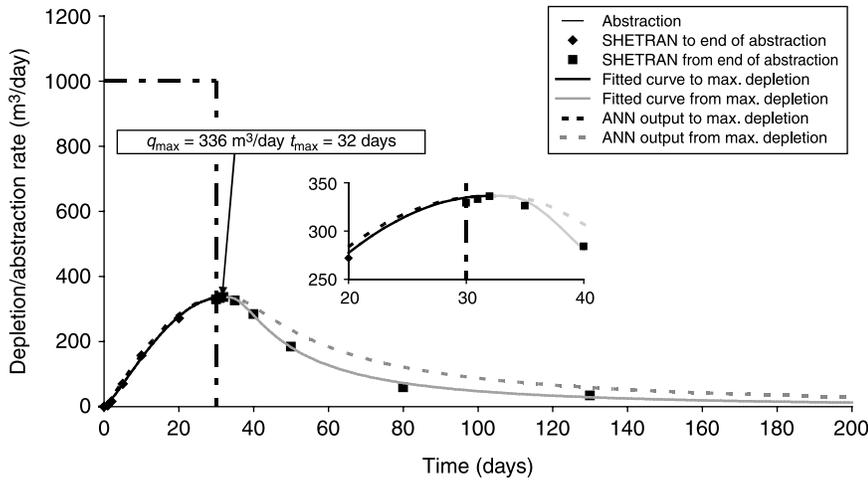| Symbol	Description	Values
D	Distance of borehole from river	212.5 m
Q	Abstraction rate	1000 m^3/d
t_{ss}	Time from to t_s to t_r (days from 0–365)	224 d
t_d	Duration of abstraction	30 d
K_v	Gravel aquifer hydraulic conductivity	25 m/d
S_y	Gravel aquifer specific yield	0.25
K_b	River bed sediment hyd. conductivity	10 m/d
R	Mean annual recharge	300 mm/yr
R_s	Recharge seasonality	0.5

Figure 4 | Flow depletion time series for the validation data set in Table 3 with SHETRAN output, fitted curve output and ANN output.

Levenberg–Marquardt fitting method (Marquardt 1963). For the second time series the SHETRAN output produces paired data of the form $(t_{2,i}\ q_{2,i})$, where $t_{2,i}$ and $q_{2,i}$ are the time and the corresponding flow depletions for the n_2 data points from the time of maximum depletion to the end of the simulation, 25 years after the end of the abstraction, and t_{end} and q_{end} are the times and flow depletion after 25 years. Logarithms of the time are taken here because a lot of the curves have a very steep drop from the maximum depletion rate and a very long tail where the depletion is zero. The normalized time fraction t_{frac} is then defined as

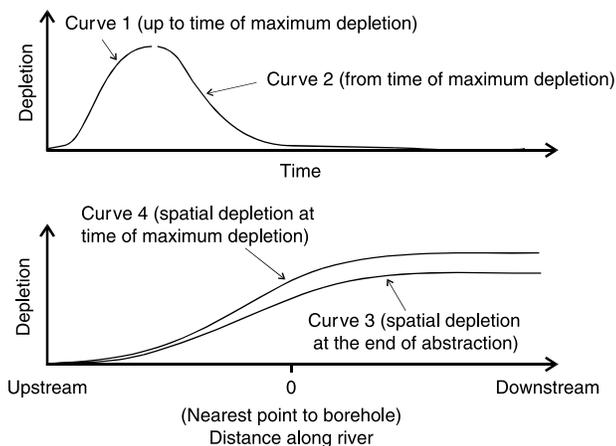$$t_{frac} = \frac{\log(t_{2,j} - t_{\max} + 1)}{\log(t_{end} - t_{\max} + 1)} \tag{3}$$

The form of the hypsometric equation for the second fitted curve is similarly

$$\frac{q_{2,j} - q_{end}}{(q_{\max} - q_{end})} = 1 - \left[1 - (1 - t_{frac})^{\frac{1}{a_2}}\right]^{(p_2 - 1)} \tag{4}$$

where a_2 and p_2 are the curve fitting parameters. The two fitted curves for the dataset in Table 3 are shown in Figure 4, which shows an excellent fit between the fitted curve and the SHETRAN output. The fitted curves for all 296 simulations can be seen in Figures 6 and 7, with the curves from the validation data in Table 3 being shown in bold. Figure 6 shows that, in many cases, there is a rapid increase in flow depletion at the start of the simulation but then the depletion stays around the maximum until the time of maximum depletion is reached. These cases generally
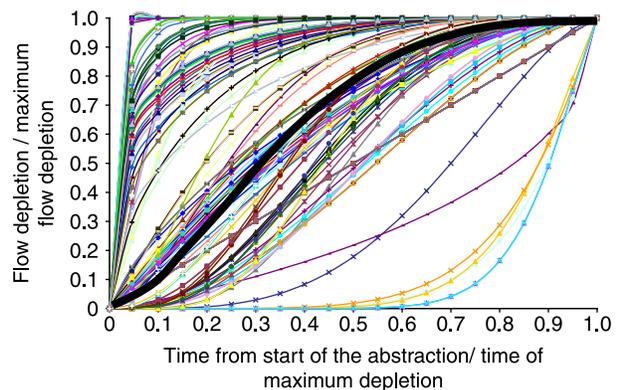

Figure 5 | Curves fitted to SHETRAN output data.

Figure 6 | Range of fitted curves for the flow depletion time series up to the time of maximum depletion. The bold curve is for the validation dataset in Table 3.
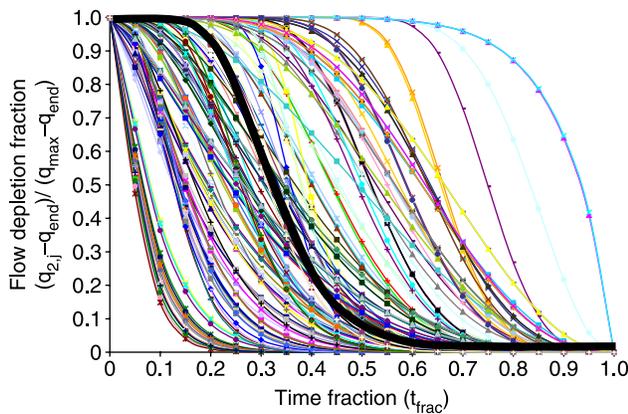
Figure 7 | Range of fitted curves for the flow depletion time series from the time of maximum depletion. Details of the terms flow depletion fraction and time fraction can be seen in Equations (3) and (4). The bold curve is for the validation data set in Table 3.

have either the abstraction well near the river and/or a high transmissivity. Another group displays an S shape, indicating that there is some time before the borehole abstraction has an effect on the river. These cases generally have a long abstraction period and the abstraction well at some distance from the river. The curves in Figure 7 for the falling part of the river depletions also show a wide range of responses. These range from those with an immediate drop and a long tail to those where the abstraction rate stays near the maximum for a long time before dropping, representing similar physical causes. The bold curves in Figures 6 and 7 correspond to the fitted curves in Figure 4 (although the bold curve in Figure 7 looks different due to the log scale on the *x* axis).

The third fitted curve was for the spatial extent of flow depletion along the river at the end of the abstraction. The extent of the depletion profile is based on the distance
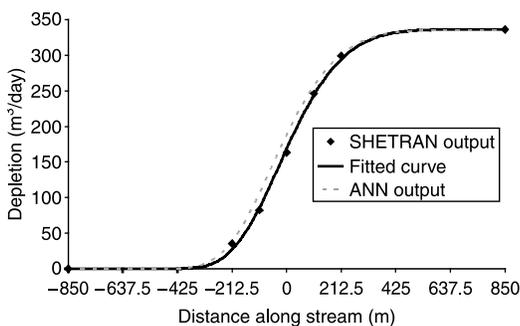

Figure 8 | Flow depletion along the river channel for the validation dataset in Table 3. Values are at the end of the abstraction period for SHETRAN output, fitted curve output and ANN output.

D from the abstraction well to the river, assuming zero depletion at a distance of 4*D* upstream and maximum depletion reached at a distance of 4*D* downstream (this extent was shown to be adequate from some preliminary studies using extreme parameter values). The value of the maximum flow depletion at the end of the abstraction is q_d. This can be calculated from the first fitted curve with parameters a_1 and p_1 at time t_d (Equation (2)). Thus

$$q_d = q_{max}\left[1 - \left(1 - \frac{t_d}{t_{max}} \right)^{\frac{1}{a_1}} \right]^{(p_1 - 1)} \tag{5}$$

The data to fit the third curve is obtained from values 25–29 (q_x) in the SHETRAN output variables (Table 2) and is of the form ($x,q_{x,d}$), where x is the distance along the river and $q_{x,d}$ is the flow depletion for these points at the end of the abstraction. The form of the hypsometric equation for the third fitted curve is then

$$\frac{q_{x,d}}{q_d} = \left[1 - \left(1 - \frac{x + 4D}{8D} \right)^{\frac{1}{a_{r1}}} \right]^{(p_{r1} - 1)} \tag{6}$$

where a_{r1} and p_{r1} are the curve fitting parameters which produce the best fit to this equation for the five pairs of points ($x,q_{x,d}$) using the Levenberg–Marquardt method.

The fitted curves for spatial depletion along the river at the end of the abstraction for the validation dataset in Table 3 are shown in Figure 8. This shows good fit between the fitted curve and the SHETRAN output. The range of fitted curves for all simulations can be seen in Figure 9. This shows two main groups, those for which there is a nearly linear increase in depletion between upstream and downstream, and those
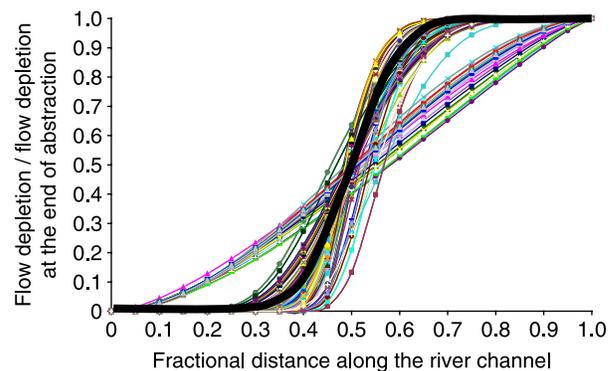

Figure 9 | Range of fitted curves for the flow depletion along the river channel at the end of the abstraction. The bold curve is for the validation data set in Table 3.

for which the depletion is restricted to a smaller region around the abstraction. The first group are produced when there is a combination of a large value for aquifer hydraulic conductivity and a long period of abstraction. In these cases the cone of depression around the well is shallow but it covers a large area and so water is depleted along a long section of the river. The second group are produced under the opposite conditions, with a low value for hydraulic conductivity and a short period of abstraction. Ideally, there should not be two distinct groups but a range of curves. However, the combination of parameter values means the flow depletion ends up in one or other of the groups, which shows there was some difficulty in covering certain parts of the response surface. A different sampling procedure in which eight values were selected for each parameter value (instead of four at the moment) but an orthogonal array with a strength of one chosen. i.e. OA(64,8,8,1), might have helped overcome this problem.

The fourth fitted curve was the depletion profile at the time of maximum depletion and is very similar to the third fitted curve. The data to fit the curve is obtained from values 30-34 (q_x) in the SHETRAN output variables (Table 2). The form of the hypsometric equation for the fourth fitted curve is then

$$\frac{q_{x,\max}}{q_{\max}} = \left[1 - \left(1 - \frac{x + 4D}{8D} \right)^{\frac{1}{a_{r2}}} \right]^{(p_{r2} - 1)} \tag{7}$$

where $q_{x,\max}$ is the flow depletion for the five data points at the time of maximum depletion, and a_{r2} and p_{r2} are the curve fitting parameters selected to produce the best fit to this equation using the Levenberg–Marquardt method. The fitted curve for the validation data set in Table 3 is nearly identical to that in Figure 8, as for this example the maximum rate of depletion occurred only two days after the end of the abstraction. The range of fitted curves is also very similar to that in Figure 9 and so is not shown.

For each of the four fitted curves the comparisons between the fitted and the measured curves from SHETRAN were excellent, producing r^2 values greater than 0.995.

ANN TRAINING

ANN training was carried out using the 9 SHETRAN input parameters in Table 1 and the 22 fitted curve output data in Table 4. Following recommended practice, the input–output data were processed before training the ANN by making the output variables dimensionless where appropriate (dividing the maximum flow depletion and the time of the maximum flow depletion by the abstraction rate and the abstraction duration, respectively), taking logarithms of the variables which spanned several orders of magnitude (duration of abstraction, the aquifer hydraulic conductivity and the bed sediment hydraulic conductivity), and normalizing all values.
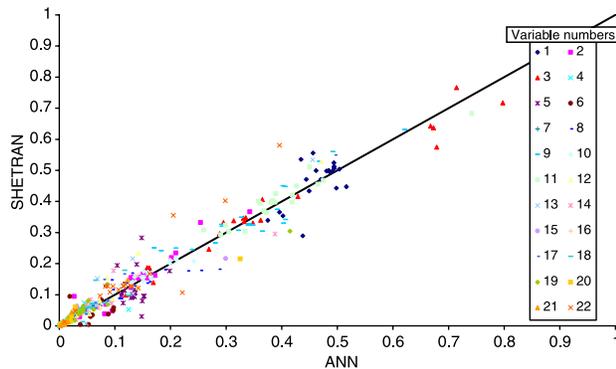
Table 4 | Curve fitted data used as output for training the ANN

Symbol	Description	Units	No. of values	Variable number
a_1	Curve shape a for flow depletion curve up to time of max. depletion	–	1	1
p_1	Curve shape p for flow depletion curve from the time of max. depletion	–	1	2
$q_{\max}/Q$	Max. flow depletion/abstraction rate	–	1	3
$t_{\max}/t_d$	Time of max. flow depletion/abstraction duration	–	1	4
a_2	Curve shape a for flow depletion curve from the time of max. depletion	–	1	5
p_2	Curve shape p for flow depletion curve from the time of max. depletion	–	1	6
q_{end}/Q	Depletion at after 25 years/abstraction rate	–	1	7
a_{r1}	Curve shape a for depletion profile in river at end of abstraction	–	1	8
p_{r1}	Curve shape p for depletion profile in river at end of abstraction	–	1	9
a_{r2}	Curve shape a for depletion profile in river at time of max. depletion	–	1	10
p_{r2}	Curve shape p for depletion profile in river at time of max. depletion	–	1	11
d	Aquifer drawdown at the end of abstraction and time of max. depletion	m	10	12–21
d_w	Drawdown in the well	m	1	22

Figure 10 | Comparison between SHETRAN output and ANN output for each of the variable numbers (see Table 4) and all of the validation datasets.

A three-layer feed-forward ANN (Beale & Jackson 1990) was then set up and trained using a back-propagation technique to minimize the root mean square error between the output from the ANN and the fitted curve output data. This was selected as it was the most popular neural network architecture currently being used in hydrology (Dawson & Wilby 2001). The ANN was trained using a back-propagation technique to minimize the root mean square error between the output from the ANN and the fitted curve output data. The ANN used 9 input nodes, with 11 hidden nodes and 22 output nodes with 256 calibration datasets and 20 test datasets used to decide when to stop training and which of the networks gives the best results. A further 20 validation datasets were used to test the quality of the ANN in which high and low values of each of the input parameters were selected. The results obtained were reasonably good with a root mean squared error of 0.051 for the calibration and 0.034
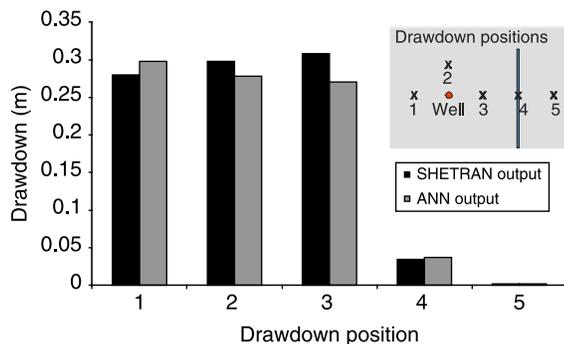
for the validation. Figure 10 shows the comparison between the SHETRAN output and ANN output from the 20 validation datasets. Overall, there is reasonably good correspondence between the two for all the output variables. The greatest errors were for variables 5 and 6 (see Table 4), the shape parameters from the time of maximum depletion.

The output data from the ANN was then used to reconstruct the complete flow depletion time series and flow depletions along the river. These data can then be compared with the curves fitted to the SHETRAN output. Figure 4 shows the output from the fitted curve and the ANN output for the validation data set in Table 3 (an example of one of the 20 validation simulations). This shows that the time of maximum depletion and the maximum depletion are very well modelled for this example. The comparison between the two curves up to this point is also excellent. From this point the ANN output drops too slowly in comparison to the fitted curve as a result of a slight error in one of the curve parameters. An excellent comparison can also be seen between the fitted curve and the ANN output for flow depletions along the river channel (Figure 8). The comparison of water table drawdowns between the SHETRAN output and the ANN output can be seen in Figure 11. The comparison is directly between the SHETRAN output and the ANN output in this case, as no curves were fitted to the data. The comparison is also excellent with similar higher values near the well, a lower value under the river and a very low value on the far side of the river.

DISCUSSION AND CONCLUSIONS

A neural network model has been used to mimic the results from numerical models of interactions between ground-water and rivers. This application of neural networks involved a large number of input parameters, each ranging over several orders of magnitude, and several types of time series and spatial output data. The ANN in this case was trained using 256 datasets, plus 20 sets used for testing during the training. An additional 20 sets were used for validation. However, with 9 input values, 11 hidden nodes and 22 output values the number of weights in this model

Figure 11 | Water table drawdown for the validation data set in Table 3. Values are at the end of the abstraction for SHETRAN output and ANN output.

was 341. A review of discussions concerning the number of training data sets and the number of weights is given in Maier & Dandy (2000). Generally, the number of data used for training an ANN should, in the case of the back-propagation algorithm, be at least as large as the number of weights (Rogers & Dowla 1994). Others state that the number of training cases should be twice (Master 1993) or ten times (Weigend *et al.* 1990) the number of model weights. However, these discussions are all based on raw data in which there is almost certainly "noise" in the input and output data and there may be no self-consistent relationship between the input and output data. In the hybrid approach described in this paper, the data are generated from a numerical model, ensuring smooth self-consistent relationships. Also, the technique (orthogonal array) that we used to select parameter values from the multi-dimensional parameter space is a very efficient design which provides better overall experimental results with less testing, as it ensures that the datasets are reasonably independent. Preliminary work showed that increasing the number of training sets produced only a marginal increase in the quality of the fitting and would have required considerable additional effort in carrying out more SHE-TRAN simulations. The efficiency of the neural network training was improved by a data reduction approach involving fitting curves to the outputs from the numerical model without significant loss of information. It is concluded that this hybrid approach, of using an artificial neural network to mimic numerical model simulations, is a feasible method of providing rapid access to the results of detailed process-based simulations, provided that appropriate data processing methods are used.

ACKNOWLEDGEMENTS

REFERENCES

Anctil, F. & Lauzon, N. 2004 Generalisation for neural networks through data sampling and training procedures, with applications to streamflow predictions. *Hydrol. Earth Syst. Sci.* **8**, 940–958.

Balkhair, K. S. 2002 Aquifer parameters determination for large diameter wells using neural network approach. *J. Hydrol.* **265**, 118–128.

Beale, R. & Jackson, T. 1990 *Neural Computing an Introduction*. Institute of Physics, Bristol.

Chua, L. H. C. & Holz, K. P. 2005 Hybrid neural network-finite element river flow model. *J. Hydraul. Engng., ASCE* **131**, 52–59.

Daliakopoulos, I. N., Coulibaly, P. & Tsanis, I. K. 2005 Groundwater level forecasting using artificial neural networks. *J. Hydrol.* **309**, 229–240.

Dawson, C. W. & Wilby, R. L. 2001 Hydrological modelling using artificial neural networks. *Prog. Phys. Geogr.* **25**, 80–108.

Dibike, Y. B. & Abbott, M. B. 1999 Application of artificial neural networks to the simulation of a two dimensional flow. *J. Hydraul. Res.* **37**, 435–446.

Ewen, J., Parkin, G. & O'Connell, P. E. 2000 SHETRAN: Distributed river basin flow and transport modeling system. *J. Hydraul. Eng., ASCE* **5**, 250–258.

Hantush, M. S. 1959 Analysis of data from pumping wells near a river. *J. Geophys. Res.* **64**, 1921–1932.

Hedayat, A. S., Sloane, N. J. A. & Stufken, J. 1999 *Orthogonal Arrays: Theory and Applications*. Springer-Verlag, Berlin.

Hunt, B. 1999 Unsteady stream depletion from groundwater pumping. *Groundwater* **37**, 98–102.

Kumar, A. R. S., Sudheer, K. P., Jain, S. K. & Agarwal, P. K. 2005 Rainfall-runoff modelling using artificial neural networks: comparison of network types. *Hydrol. Process.* **19**, 1277–1291.

Maier, H. R. & Dandy, G. C. 2000 Neural networks for the prediction and forecasting of water resources variables: a review of modelling issues and applications. *Environ. Modell. Softw.* **15**, 101–124.

Marquardt, D. 1963 An algorithm for least-squares estimation of nonlinear parameters. *J. Appl. Math.* **11**, 431–441.

Master, T. 1993 *Practical Neural Network Recipes in C++*. Academic Press, New York.

Ochoa-Rivera, J. C., Garcia-Bartual, R. & Andreu, J. 2002 Multivariate synthetic streamflow generation using a hybrid model based on artificial neural networks. *Hydrol. Earth Syst. Sci.* **6**, 641–654.

Parkin, G., Birkinshaw, S. J., Younger, P. L. & Rao, Z. 2007 A numerical modelling and neural network approach to estimate the impact of groundwater abstractions on river flows. *J. Hydrol.* **339**, 15–28.

Rao, Z. & Jamieson, D. G. 1997 The use of neural networks and genetic algorithms for design of groundwater remediation schemes. *Hydrol. Earth Syst. Sci.* **1**, 345–356.

Rao, Z. & O'Connell, P.E. 1999 Integrating ANNs and process-based models for water quality modelling. In *Proceedings of*

the Second Inter-regional Conference on Environment-Water, 1–4 September, Lausanne, Switzerland (ed. A. Musy, L. S. Pereira & M. Fritsch), PPUR, EPFL, Lausanne.

Rogers, L. L. & Dowla, F. U. 1994 Optimization of groundwater remediation using artificial neural networks with parallel solute transport modeling. *Wat. Res. Res.* **30**, 457–481.

Shigidi, A. & Garcia, L. A. 2003 Parameter estimation in groundwater hydrology sing artificial neural networks. *J. Comput. Civ. Engng.* **17**, 281–289.

Theis, C. V. 1941 The effect of a well on the flow of a nearby stream. *Am. Geophys. Union Trans.* **22**, 734–738.

Tokar, S. A. & Johnson, P. A. 1999 Rainfall-runoff modeling using artifical neural networks. *J. Hydrol. Engng.* **4**, 232–239.

Weigend, A. S., Rumelhart, D. E. & Huberman, B. A. 1990 Predicting the future: a connectionist approach. *Int. J. Neural Syst.* **1**, 193–209.

Winter, T. C. 1995 Recent advances in understanding the interaction of groundwater and surface-water. *Rev. Geophys.* **33**, 985–994.

First received 9 February 2007; accepted in revised form 26 August 2007