7

Real Neural Storage Radiator Control

7.1 Background

The simulation results from the previous chapter indicate that neural networks could be used to improve storage radiator control. Computer simulations can significantly reduce the development time of new products, but more practical problems may only come to light once working prototypes are developed. The development and results from a prototype of what is believed to be the world's first neural network controlled domestic storage radiator are reported in this chapter.

The room being controlled was on the first floor of an occupied two-bedroom property located close to Bedford town centre. A plan of the first floor and photograph of the exterior are shown in Fig 7-1 and Fig 7-2.

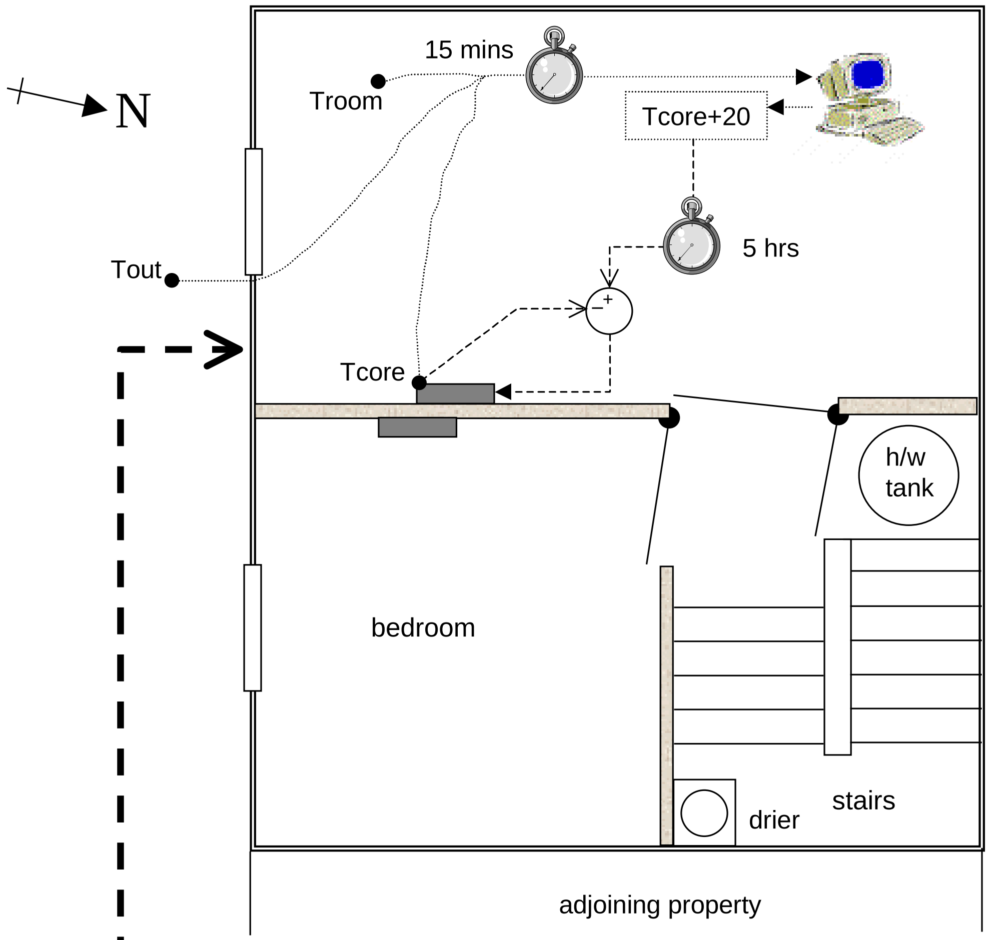

Fig 7-1 A *Plan of the first floor and schematic of the control scheme*

Fig 7-2 *A photograph of the exterior of the house identifying the room under control*

The house is a converted nursery that was constructed around the turn of the century. The walls are solid 9-inch thick brick and there is a large glass skylight above the stairs. The room being monitored was the spare bedroom which was used as a study and had a 1.7kW storage radiator initially on E7, as did the bedroom. Below the study is a living/kitchen area with a storage-fan-convector heater on E7. Below the bedroom is the bathroom. On the landing of the stairs was a clothes drier that was used from time to time with and without an extraction hose venting through the skylight.

Temperatures logged were ambient, room (×2), storage heater core, storage heater air outlet and a desk lamp temperature. This lamp temperature was included as a potential means of identifying occupancy and thus heat gains from computer equipment, lighting and body heat. Logging commenced on 25th September 1997 and continued in 15-minute intervals until 14th May 1998. The neural controller was in operation from the 25th February with the period 17th – 30th March being the final ‘de-bugged’ version.

Fig 7-3 shows the ambient, core and room temperatures for November, when the heater was being charged on an E7 schedule. It can be seen how adjusting the input charge setting changes the maximum allowable core temperature. The reliance of the room temperature on ambient is evident, with there being little time delay between the underlying room temperature and ambient. This would suggest that the house is relatively lightweight. The responsiveness to the core temperature can also be seen, with the maximum room temperature occurring at 7:00 am when the core is hottest, confirming how the existing heater does not give the heat when it is required.

Fig 7-4 shows the whole heating season. The charge was initially set on E7 and then varied by disconnecting it from the off-peak circuit and controlling it by means of a timer plug. This change was introduced so that a diverse range of data could be collected. Several periods were included when the heater was switched off and the core allowed to drop to room temperature. This gave data at the lower extreme and an indication of the rate of heat loss of the core at different temperatures. A second reason for switching off the heater can be seen by the relatively high room temperature that was being maintained.

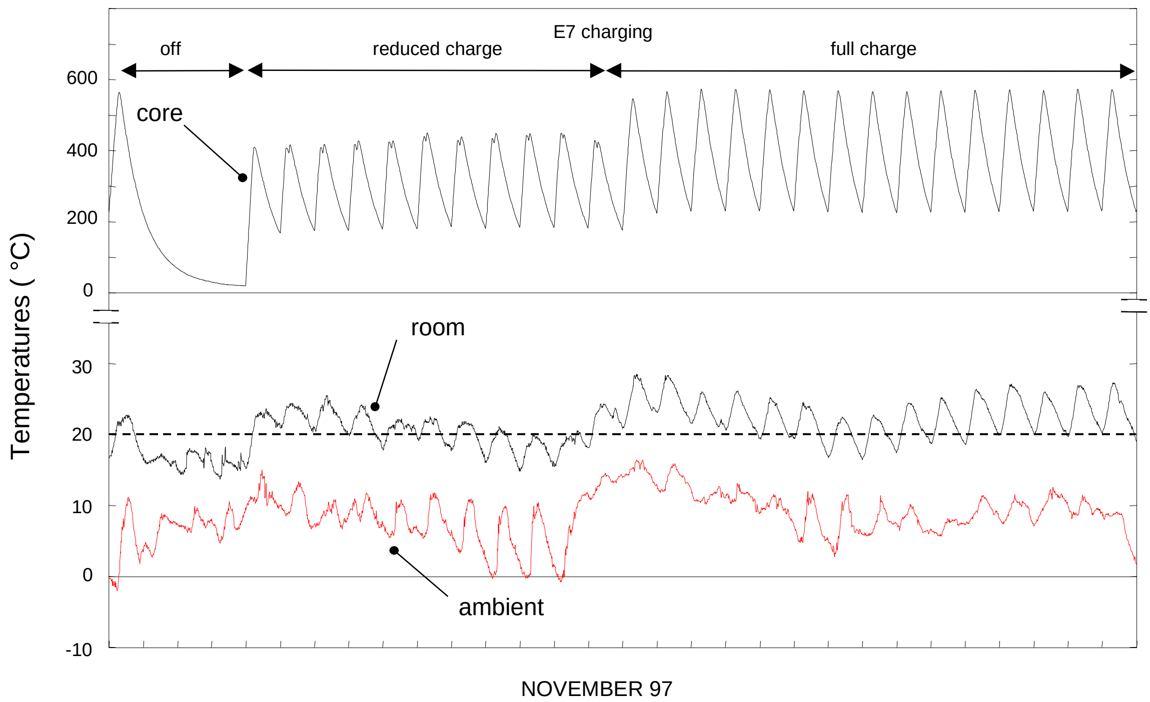

Fig 7-3 *Room, core and ambient temperatures for November*

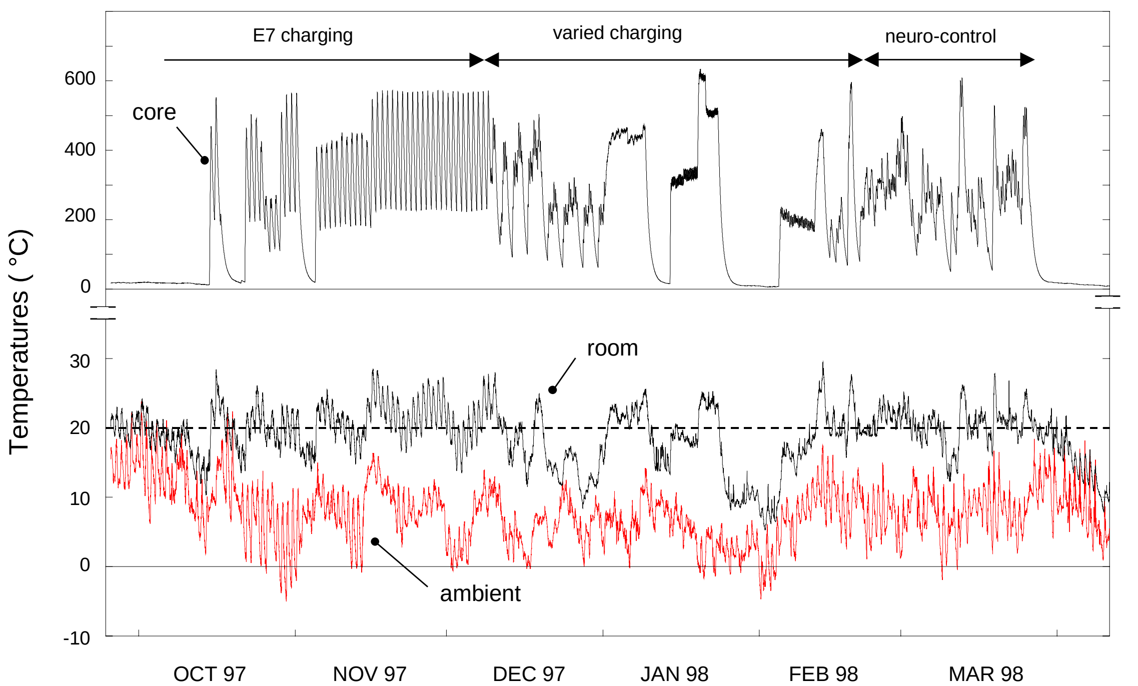

Fig 7-4 *Room, core and ambient temperatures for the whole heating season*

The heating season in which the work was undertaken was unusually warm. Because it was the last winter of the project, in order to gather data, the heating had to be on when normally it would not have been. The following are press releases from the Met. Office describing the weather at the time,

ENGLAND 1997: TURNED OUT WARM AGAIN - This year has been the third warmest in England since records began over 300 years ago. August in 1997 was the second warmest on record, and February and March were both particularly warm months over England.

SCORCHING START TO 1998 - The first six months of 1998 have been easily the warmest first half of a year globally since reliable records began in 1860. Provisional observations analysed jointly by The Met. Office and the University of East Anglia show that the temperature averaged over January - June 1998 has been some 0.6°C greater than the average climate (calculated from the period 1961-1990). Each individual month in 1998 so far has been the warmest such month on record. Furthermore, the twelve month period July 1997 to June 1998 (with an anomaly of 0.56°C) has been very much warmer than any other 12 month period not influenced by the current El Niño.

Not surprisingly, by tomorrow (Friday) morning April 1998 will be the wettest April this century.

7.2 Data Analysis

Data from September to February was analysed in order to create a fixed model that would be placed in the controller.

The simulations in the previous chapter used electrical charge as the control variable. In the practical tests only temperatures were measured, so the control decision is to predetermine a future core temperature that will track a future set point temperature. The control variable is thus the core temperature.

The control horizon in the simulations was up to 24 hours and depended on knowing future ambient temperatures, which were used retrospectively. In reality a weather forecast is required, the only option available for the prototype being an ‘educated guess’ (using a neural forecast), which limited the prediction horizon that could be used.

The importance of ambient temperature to the system determines how far ahead it would be reasonable to model, knowing that ambient predictions are likely to be poor. Several neural networks were trained to model a future room temperature given the current core, room and ambient temperatures. Both the intervening core and ambient temperatures were included and then only the intervening core temperatures. This was done in order to quantify the importance of the ambient temperature. In the neural networks it was found that one neuron was sufficient to model the system. A naïve prediction, which simply says that the future temperature will be the same as the current temperature, was also made. This is the worst case model and should always be used for evaluating neural models to ensure that they are doing more than simply repeating the last known value. The results are shown in Fig 7-5.

The intervening core and ambient temperatures determine the major heat transfer processes into the room and will thus influence the future room temperature. By including both of these in the model the rms error for a 15 hour horizon was 1.07 °C, compared

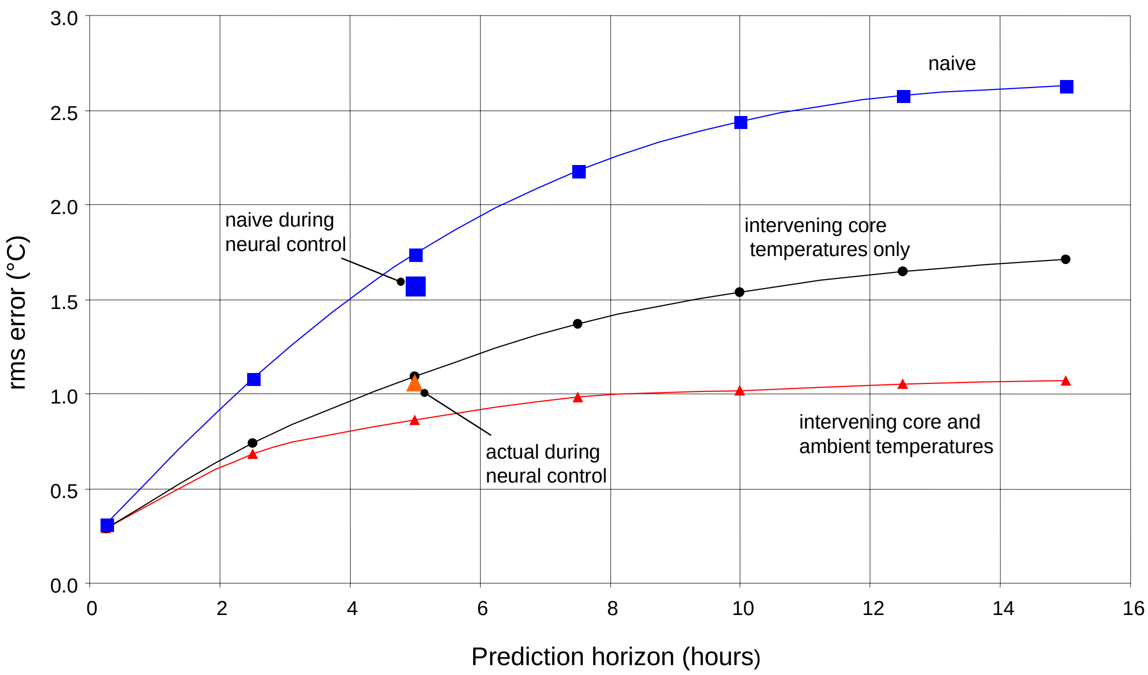

Fig 7-5 *Neural model errors for predicting future room temperatures. Inputs are the current core, room and ambient temperatures and intervening ambient and/or core temperatures. A naïve prediction assumes that the future room temperature will be the same as the current. The data used is for the period September-February.*

with 1.71 °C by only including the core temperatures and 2.63 °C for a naïve prediction. The longer the prediction horizon the more important it becomes to know the ambient temperatures, although just knowing the core temperatures is still a significant improvement over a naïve prediction. These results show that a network with a single hidden neuron gives a reasonable model and suggest that linear regression might work just as well. A prediction horizon of 5 hours was chosen, for which the difference in rms errors between knowing and not knowing the intervening ambient temperatures was 0.23°C for the data set analysed.

If a building is left for long enough then the room temperature will reach thermal equilibrium with its surroundings. It is thus not necessary to know the current room temperature if a long enough prediction horizon is being used. Two models were created, one included the previous 20 hours core and ambient temperatures but not the initial temperature 20 hours ago, while the other included the past 5 hours core and ambient temperatures but also the initial room temperature, as was used in the controller. The errors of these models are shown in Fig 7-6.

The rms error for model 1 is at a level that is only marginally better than a naïve prediction. This might have been be discarded as poor but a close analysis of the errors gives some meaningful information.

The model indicates that something different is occurring at the end of October and December, shown by the constant bias from zero error. The last few days in October saw the first cold front, as can be seen from the ambient temperature of Fig 7-4. The downstairs heater had a controller that would keep the room at a set temperature by activating either the fan to release stored heat or the convector, which was available 24-hours. During this period this control was set so that there was a constant heat output throughout the day in order to keep off the chill. Normally this was only activated during occupancy. The effect of this on the room above is that there is an increased level of heat input through the floor, which changes the model and is why the errors are seen to be constantly underestimating the room temperature. This extra heat input to the room can be seen to equate to 3-4 °C.

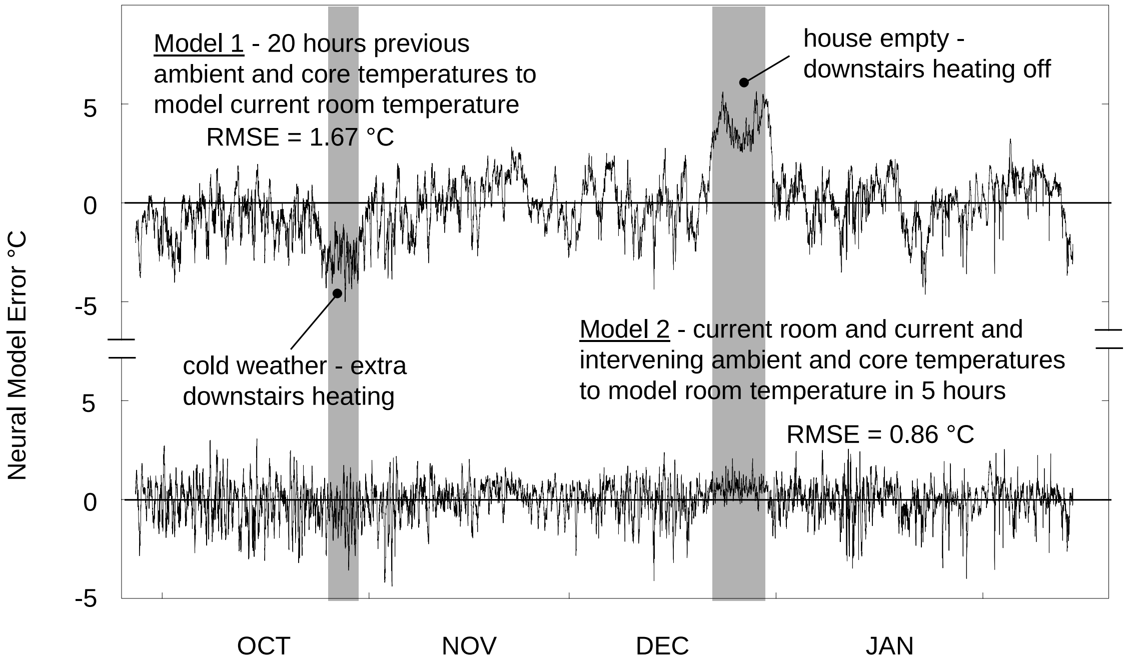

Fig 7-6 *Not including the current temperature in the model gives a worse error but anomalies are easier to spot*

At the end of December the house was unoccupied and the downstairs heating switched off. Model 1 clearly identifies the change by showing that the room is about 4 °C cooler than is expected, the missing heat input from the room below. Model 2 does pick up some change during this period with the errors being constantly positive, but what is more noticeable is the reduced noise level in the errors. This is due to the house being empty and identifies that there will be noise in the data caused by the occupant's behaviour.

The information extracted parallels with the causal model created for electrical load modelling in chapter 2. Including the existing temperature has the same effect as including previous loads, giving a better model by introducing a good initial guess. This can give a good model, but as has been demonstrated, causal models can reveal a lot of information even though the errors are worse. What they do is give clues as to how the model could be improved, which in this case would be to include as inputs the temperature of the room below and all other events that cause heat fluctuations, or 'noise'.

Examples of temperature fluctuations in the test room that are caused by events other than heater or ambient temperatures are shown in Fig 7-7. The 11 day period shown includes 3 days when the house was unoccupied, seen to be a time when the room temperature is well behaved. Occupancy can be seen to introduce a certain amount of noise, with the periods when the drier was on being clearly identified as spikes.

These ‘spikes’ are short-term energy inputs that do not affect the background temperature in the long run, and as such are considered to be noise. If the current temperature happened to be during one of these brief anomalies then errors would be introduced in the model as this would not represent the ‘real’ background temperature. In an attempt to filter out this noise and give a more stable longer term temperature, the three preceding as well as the current room temperature were included in the model. From this an effective time averaged current temperature will be calculated, with the resulting model showing a reduced level of noise in the errors.

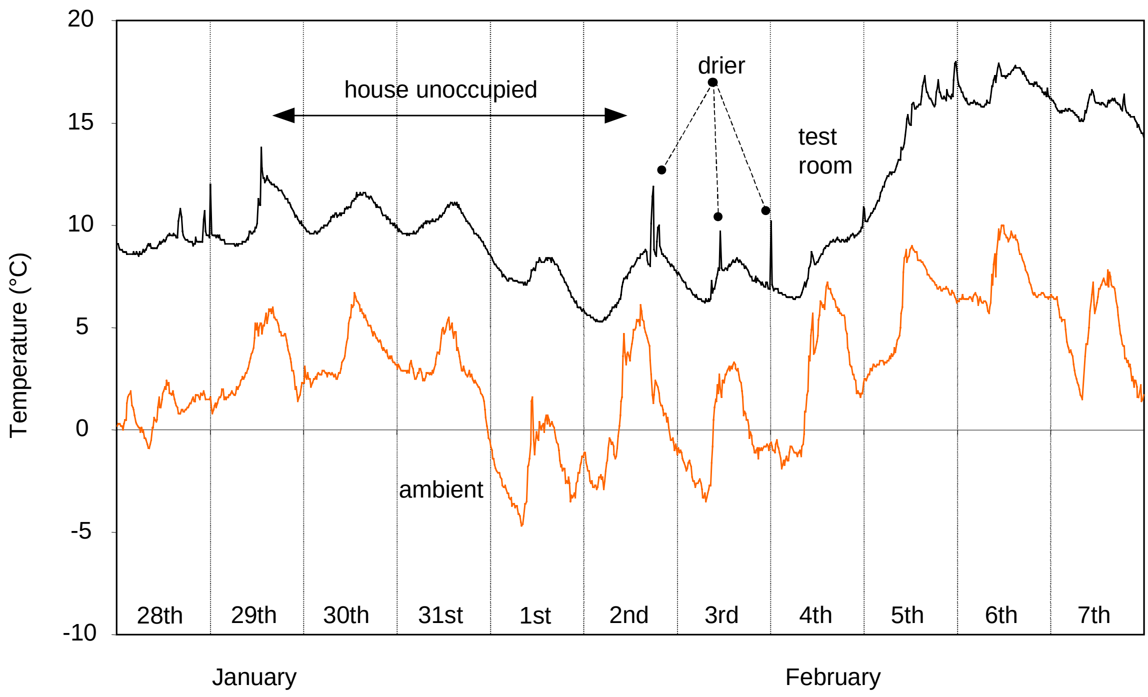

Fig 7-7 *An example of occupancy creating ‘noise’ that the network does not have causal information to model. The unoccupied period can be seen to be ‘noise’ free.*

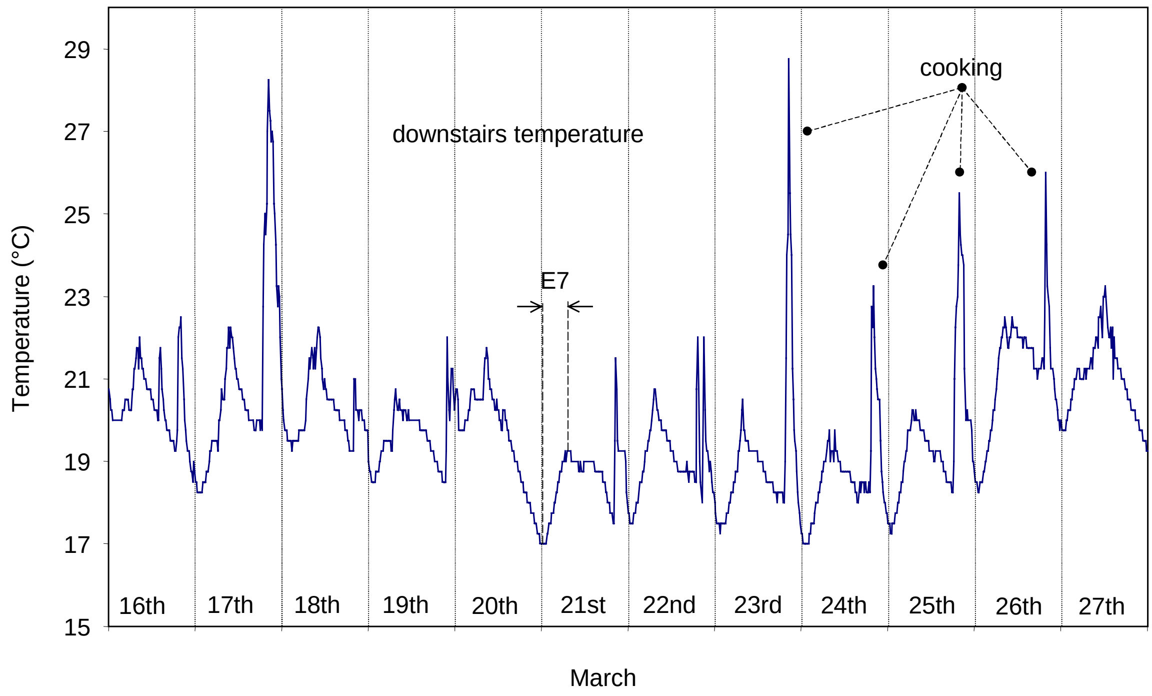

Fig 7-8 *An example of temperature fluctuations in the downstairs room*

Fig 7-8 shows an example of the temperature fluctuations that occur in the room below that will be transmitted into the test room through the floor and by natural air circulation. Short-term temperature increases in the late evening due to cooking and heater fan activation are very evident as is the underlying E7 room heating. These fluctuations will affect the test room temperature and are reasons why the noise is present in the model errors.

7.3 Neural Controller

In the prototype controller it was considered ambitious to attempt both predictive control and optimisation in the first instance. The only objective required was to track a given temperature by implementing control set points of the core temperature determined 5 hours previous. This was implemented by using the neural emulator as shown in Fig 7-9 and Fig 7-1.

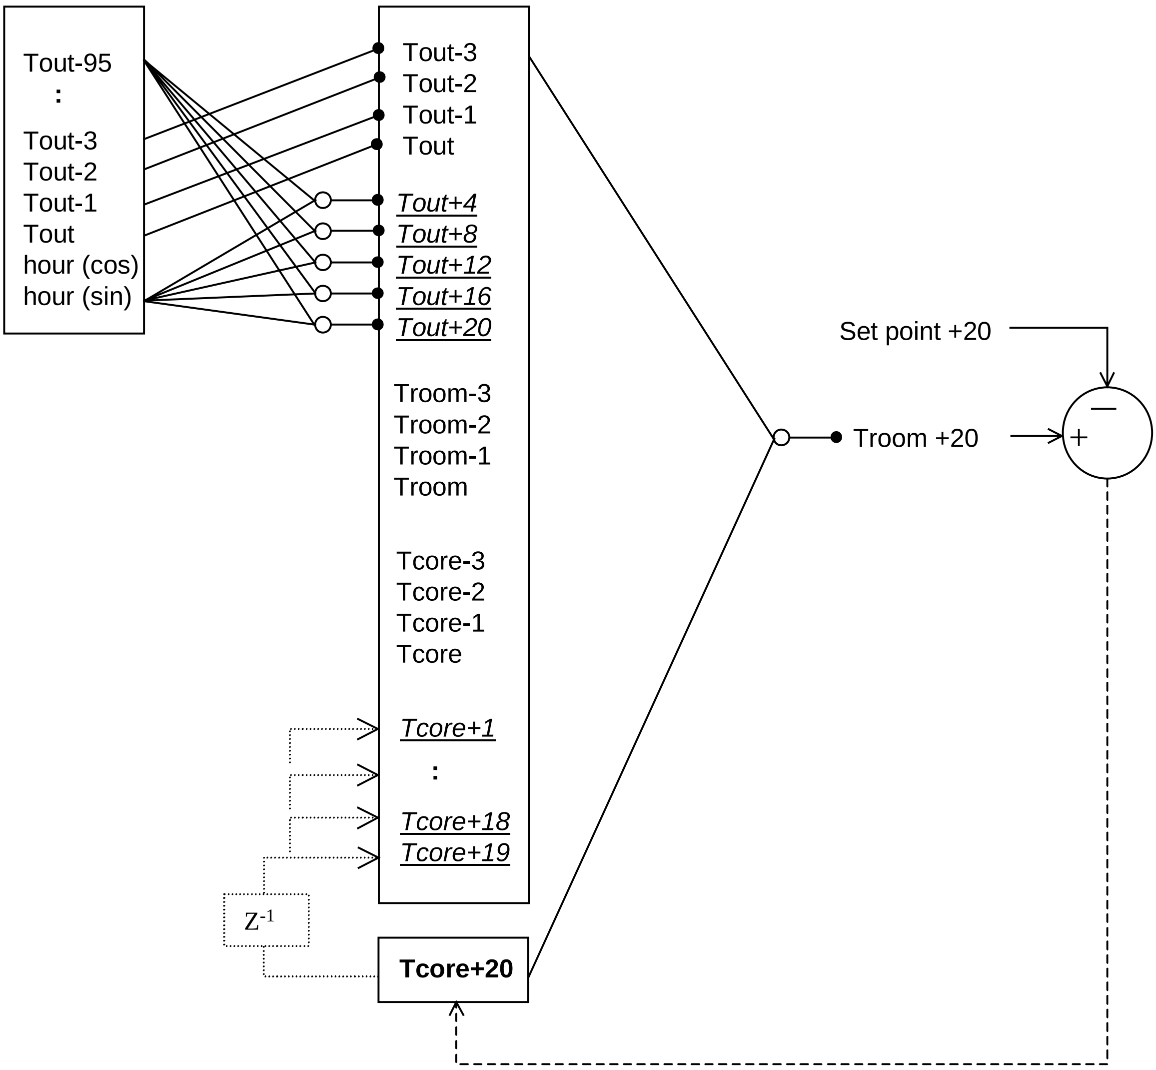

Fig 7-9 *The neural emulator created for the heater control. Every 15 minutes a new core temperature for 5 hours time is calculated (Tcore+20) by finding such a temperature that minimises the error between the set point and the model prediction. There are upper and lower bounds on this core temperature determined by what it is physically possible to achieve from its previous state. Once a core temperature is determined it is added to a stack and the controller acts as a thermostat switch that tries to achieve the core temperature at the top of the stack (Tcore+1). Once a core temperature is determined and enters the stack it cannot be adjusted.*

The emulator consisted of 6 neural networks. Five of these were to predict the ambient temperature 1,2,3,4 and 5 hours ahead. These predictions were then used by the main network to determine the required core temperature.

The network inputs to predict the ambient temperature were the previous 24 hours temperatures along with the hour of the day. The output was the temperature either 1,2,3,4 or 5 hours ahead. Including the hour of the day, coded as a sine and cosine, improved the model performance by giving some reference as to when the turning points would occur. Each network weights were trained individually, not as a single network with 5 outputs.

The network weights were calculated off line and then implemented in the controller with no further training. Fig 7-10 shows the errors of the created emulator on the training data, showing the created model predicting a temperature within 1°C in 81% of cases.

The controller hardware was an IAC600 donated by Satchwell Controls with the software custom written Visual Basic. As different temperature sensors to the data logger were used some calibration was required which is a possible source of error.

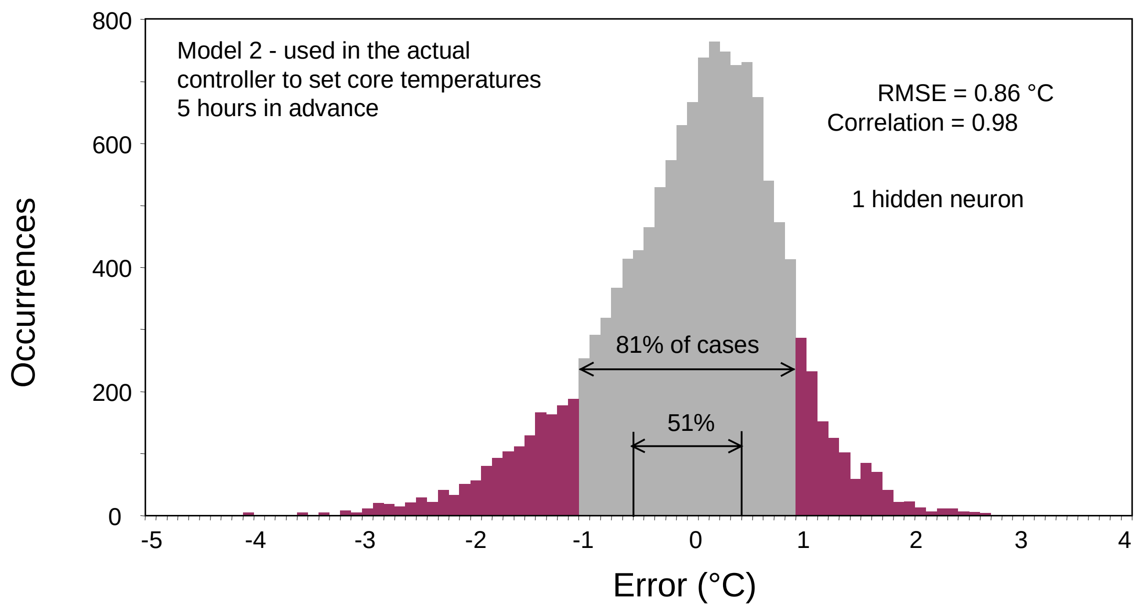

Fig 7-10 *The errors on the training data for the model implemented in the controller*

Fig 7-11 shows the results over two periods in March, indicating how the set point was continuously varied. Because storage radiators have a slow response the set point cannot always be achieved. The reported error is thus the difference between the model prediction and the achieved temperatures.

The rms error over the final two week period 17th-30th March was 1.06 °C with a naïve prediction error being 1.57 °C. How these relate to the training data is shown in Fig 7-5 (on page 151). Because there is set point temperature control the naïve prediction is better than on the training data when there was no temperature control. This is because temperatures are more constant and thus likely to be the same 5 hours ahead. The actual rms error is slightly less than would have been expected if no intervening ambient temperatures were used. This would suggest that the temperature predictions were poor, as was later confirmed.

The errors can be seen to be constantly underestimated, which is probably due to a slight calibration error. If this was not the case and they had a mean error of zero then the rms error would be significantly reduced.

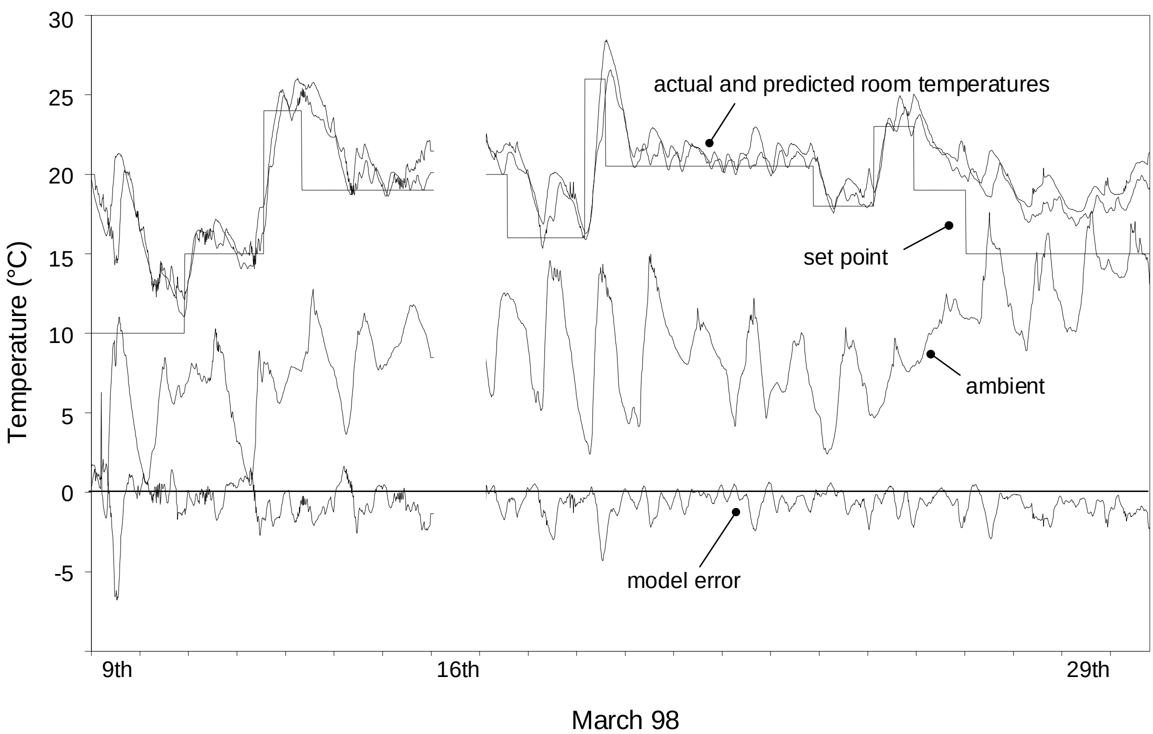

Fig 7-11 *The results of the neuro-controlled storage radiator. The 'de-bugged' controller operated constantly for two weeks from the 17th March.*

Fig 7-12 shows the bedroom temperature over the same period as for the neuro-control. The room is smaller but has an identical heater operating on E7 and is thus useful to gauge what the uncontrolled test room response might have been. Daily temperature swings of up to 4°C are common, which would have been even greater in a larger room with more exterior wall area. It is also noticeable that the room is at its coldest in the late evening when warmth will be desired. The control achieved by the neuro-controller can thus be considered a success.

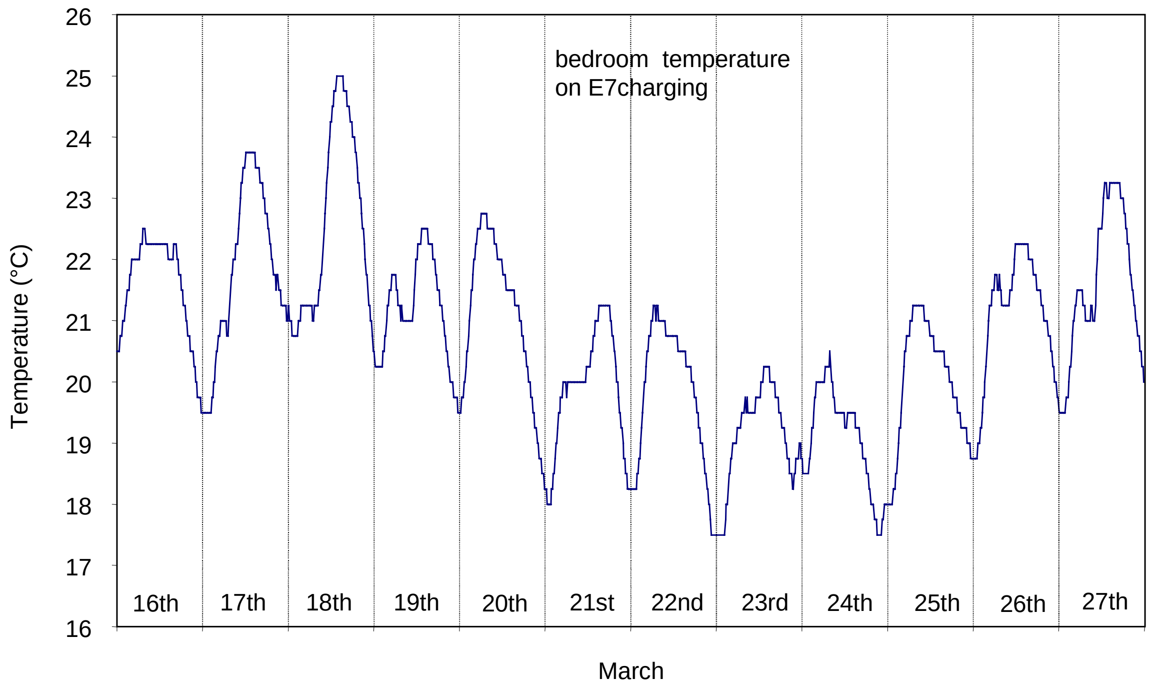

Fig 7-12 *The bedroom temperatures over the same period as for the neuro-control with an identical heater operating on an E7 charging schedule*

7.4 Chapter Summary

This chapter has reported the results of a prototype neuro-controller for a storage radiator. Important lessons have been learnt that would not be evident in simulation, namely the introduction of noise caused by occupants and the importance of heat gains through floors.

8

Project Overview

8.1 About this Chapter

The New Product Development Department of Eastern Electricity sponsored this research project to develop an intelligent controller for storage radiators. Soon after commencement the decision was taken to close the department, ending any interest they had in the research and in any future development of a product.

This chapter summarises the key stages of the work, giving recommendations of how the results might be of use to Eastern. The most important offering that can probably be given is a description of how this neural network project evolved over the four years, highlighting important lessons that were learnt. This will hopefully be of use to others embarking on similar work.

8.2 Load Forecasting

The load forecasting techniques developed should be of interest to Eastern as electricity trading in the free market will require accurate predictions to cut costs and stay ahead of the competition. Initially, all Eastern wanted to know from their data was a growth trend over the years, which was achieved, with much more besides.

Compared with existing commercial products the techniques developed deal with growth and holiday periods in much more detail. The reported results, although not directly comparable, appear superior.

American utilities are investing heavily in neural forecasting techniques, with investments in commercial software of up to $100,000 [**98**] resulting in improved forecasting accuracy [**19,20**] and financial savings. With such large amounts of money being paid for a technology in its infancy there is potential for developing commercial software, which would be very easy to accomplish.

8.3 Water Optimisation

The results from the simulations showed that savings were made by avoiding the midnight price peaks, caused by the E7 surge. This can be achieved by utilising the radio tele-switch control of off-peak heating systems. This option has not been utilised by Eastern until very recently, with switching occurring at the same times every day. Although it is in the tariff contract that times can be varied, the reasons for not exercising this control option appears to be a fear of upsetting customers. If this is the case then it was an expensive investment installing the tele-switches in the first place. Complaints from customers since the switching times were varied have been very few.

8.4 Intelligent Heating Control

Development of a commercial product for domestic storage radiator control is not an immediate necessity, as other technologies have to fall in place to provide the infrastructure required. These include half-hourly meters, real time tariffs and communication channels for weather and pool price data. The last four years has seen exponential growth in telecommunication and microchip technologies including the internet and the use of home PC's. What was probably considered a fantasy product four years ago is now very much achievable at ever decreasing costs.

It would be advisable for anyone wanting to develop the idea to start in commercial buildings, where the potential of large savings might encourage the initial investment costs. Any heating system could be optimised by learning the building characteristics, it does not have to be limited to storage radiators.

It is only a matter of time before the 'smart house' becomes a reality and there is the opportunity for money to be made.

8.5 Experiences of Pursuing a Neural Network Project

At the start of this research it has to be said that neither I, or anyone involved with the project really new what a neural network was, apart from it was something that could learn things. As a result, there was a lot of time wasted and things were probably done the hard way. The following is a list of key points that directly influenced the research.

1) A lot of early time was wasted trying to read 'introductory' textbooks and journal papers. The best and quickest way to find out what neural networks are is to have it explained.

2) During a visit to the DTI Neural Awareness Campaign I was informed that there was an abundance of ‘free’ neural simulators on the internet. One was downloaded (SNNS) and basically played with for a while.

3) A chance meeting with a fellow student who was having problems coding her own neural network (and seeking my advice!) introduced me to the idea that neural networks were not very difficult to create.

4) The development of my own code in a suitable language. Fortran 90 was chosen because it is simple to follow and has many intrinsic matrix multiplication functions, which is all neural networks are.

5) The decision to switch from a UNIX system to PC’s made the coding much more portable and it could be done at home.

6) Realising that to simulate all the various components of a controller, the whole system would have to be coded, unless a tool like MATLAB was used. It was decided that using custom code is much better for research purposes as you can control exactly what is happening and are not reliant on someone else’s algorithms. Initially a building thermal analysis package was used but it was soon realised that this could not be interfaced with a controller.

7) It was soon evident that no real understanding of how neural networks really work could be achieved by using simulated data, as they contained no noise. The event that redirected the whole research focus was the analysis of the ‘real’ electrical load data.

8) A neural model was created for the load data and presented to Eastern. It was then that a quite reasonable question was asked, ‘how do your network weights relate to my linear regression weights’. This was not the type of question people are supposed to ask as a neural network is a black box, isn’t it? It was then realised that there must be some meaning in the weights, only no one had really bothered investigating this before.

9) It was then realised that neural networks are non-linear regression and definitely NOT artificial intelligence.

8.6 Cost Analysis

Because this research has been mainly computer based the equipment costs have been minimal. The ‘extra unavoidable’ costs required were £300 for a Fortran90 compiler for PC’s and roughly £200 for wiring and thermocouples for the controller. This works out at around £2.50 per week. Neural network projects can be very cheap, the major investment required is in human resources.

8.7 The Future

Once the hard work of the first application is overcome, further applications can be achieved almost instantly. The possibilities are then just a matter of a creative imagination and inquisitive mind.

---

[98] **Ben Hobbs – International One Day Symposium on Electricity Load Forecasting – London Business School, 6th July 1998.**