


Artificial Neural Network – Iris Data

This program is a simple example of backpropagation in an ANN learning the Iris data set.



The display is very simple: it shows the current error, expressed as the number of samples that the network is not recognizing correctly, and the best ever error during the current run.

The entire group of 150 samples is treated as a single training set. The data are fed in through the code

```
double iris[][] = {
{5.1, 3.5, 1.4, 0.2, 1, 0, 0}, //0
{4.9, 3, 1.4, 0.2, 1, 0, 0},
{4.7, 3.2, 1.3, 0.2, 1, 0, 0},
{4.6, 3.1, 1.5, 0.2, 1, 0, 0},
{ 5, 3.6, 1.4, 0.2, 1, 0, 0},
{5.4, 3.9, 1.7, 0.4, 1, 0, 0},
.....
}
```

in which each line is data for a single type of Iris. The different types are distinguished by the 5th, 6th and 7th entries, exactly one of which is non-zero for each type.

BUTTONS

Buttons in the window created by the program have the following functions:

Button	Function
Play	Restart the program if it has previously been paused.
Reset	Restart the calculation from scratch; all parameters are set to their original values.
Pause	Temporarily suspend execution.
Step	Execute a single epoch.

SLIDERS

The sliders that can be adjusted by the user are:

Slider	Default	Comment
Eta	0.08	See text for an explanation of this parameter.
Alpha	1.00	See text for an explanation of this parameter.

Investigations and Exercises

The behavior of an ANN may depend in a complex way upon the size and geometry of the network, the values of the variable parameters, whether input data are scaled, the type of squashing function used, and the size and complexity of the data set. To gain an insight into how the various parameters affect the learning of a network, it is best to construct an ANN of your own so you can test this behavior.

The exercises below illustrate how a very simple network can learn.

1 Default parameters

Start the simulation and note the behavior of the error as the simulation runs. Typically, the number of misclassifications remains close to 100 for an initial period of up to a couple of hundred epochs. (The number of misclassifications per epoch should be around 100 if the network selects an Iris type at random: there are 150 samples and three types of flower, so if the network guesses randomly, it will be incorrect two times in every three, equating to 100 failures).

Run the simulation several times to see how reproducible this behavior is. Notice that there is usually an induction period, during which the ANN learns little, but that once it starts to learn the number of misclassifications falls rapidly to a low value.

2. Dependence on parameters

The behavior of the network on the parameters ε and α is not trivial. However, you will gain some insight by starting a run with a value of these parameters different from the default values and following progress of the error. Change the value of the parameters part way through a run; how is learning affected?