

The Leopard

You will be given the image of a leopard that is 300×500 pixels. Your goal: Break up the image into 5×5 non-overlapping samples. For each small sub-image, classify it as belonging to the leopard or to the background.

The work in splitting up the data is already done for you in the sample script, all you need to do is set up and train the neural network. The script sets up the data X as 25×751 (751 sample points for training) and **Targets** is 1×751 .

First, use the default values for training, and train several times to see if there is much of a difference. You may assume that you need to use 10 nodes in the hidden layer.

When you're finished training, have the network output the values for all the targets and plot them (the original data is in order).

When you're finished with the default methods, try using regular gradient descent. To do this, you will need to change some network values (in the script, do this after the `newff` command):

```
net.trainFcn='traingd';  
net.trainParam.epochs=1500;
```

For extra practice, be sure you can change the default percentages used in training, testing and validation (See page 5 of the notes on backpropagation of error).

Comment on your results.




Figure 1: The Leopard