Machine Vision

Veneer Grading with Artificial Neural Networks

Ventek, Inc.
4030 West 1st Ave., STE 100
Eugene, OR 97402

Phone: 1-541-344-5578

Author: James Gibbons, Software Engineer and GS2000 Architect.

Presented at Scan Pro 2000 by Rodger Van Voorhis, President of Ventek.

Human Vision Process

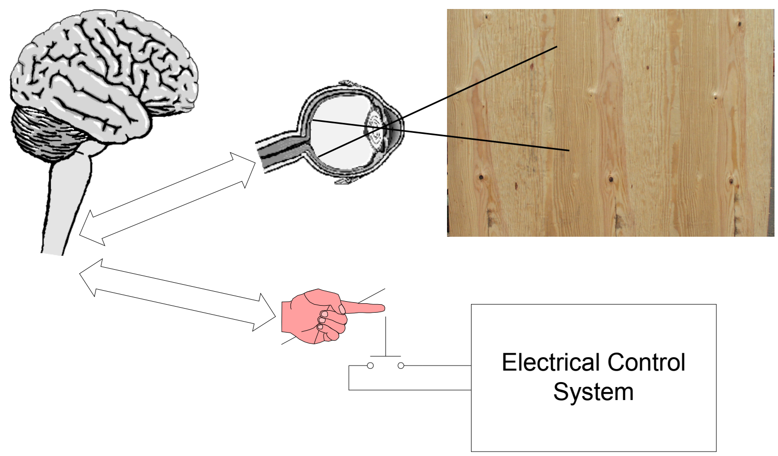

Object-Eye-Brain Interaction

- Eye views object.
- Eye sends processed information to brain.
- Brain directs eye to scan object.
- Repeat as necessary or until time runs out.

Human vision is a scanning process. The brain directs the eye to scan over an object and select areas of interest. Once an area is identified and processed, the brain moves the eye to other areas.

The eye doesn't send a simple copy of the image to the brain. Instead there are layers of nerve cells between the rod and cone light sensing cells and the optic nerve. These nerve cells pre-process the visual information into a higher level than simple intensity levels, although intensity is part of the information sent to the brain.

Brain-Hand Interaction

- Brain builds model of object.
- Brain compares model to prior experience to determine a grade.
- Brain directs hand to push button.
- Control system uses input to process object.

During the eye-object scanning process, the brain is building up a model of the object. As more of the object is scanned, the model is refined to higher degrees of accuracy. Exactly how this is done is still a very hot area of research.

Human Vision Process

- Eye pre-processes image before sending it to brain.
- Eye only sees a small area with detail.
- Brain is a massively parallel system and is poorly understood at this point in time.
- Brain must be trained for many years before it can model objects, make decisions, and act on them.

Only a small area of the retina (called the *fovea centralis*) is sensitive to color and detail. This is why the brain must direct the eye to scan a large object and key into areas of interest before it can build a model of the object.

While many details of the brain's visual processing are well understood, the full interaction of the visual modeling and decision process are poorly understood at this time. The visual modeling actually involves a full 3-dimensional processing with corrections for lighting, shadow and color effects. Studies of brain damaged subjects have shown that various aspects of these processing abilities are present in specific brain areas because they can be lost if the areas are damaged.

Machine Vision Process

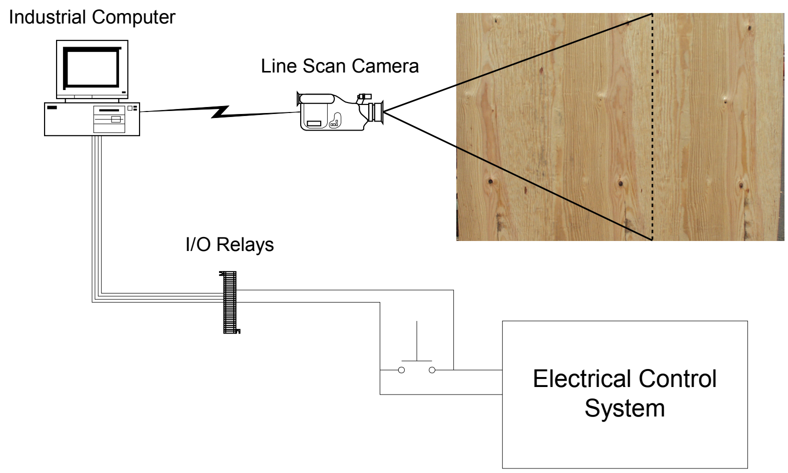

Camera Process

- Camera captures picture elements (pixels) and converts them to digital values (light intensity numbers).
- Pixels are sent in line or matrix format to image processor.

Smart cameras have been developed that can pre-process the image information before sending it to the main vision processor. In most cases, a small computer chip is added to a standard camera and it performs some specific processing before the data is sent to the main processor. It is generally better to keep the camera simple and perform all the processing at a single location.

Other smart cameras are adding the processing elements onto the same chip as the light sensitive elements. In this case the processing can be massively parallel with several operations run in parallel with the light sensing. This type of camera is still in the experimental stage, but some commercial versions exist.

The range of light levels that a silicon CCD camera can sense is limited compared to the human eye. The eye is very good at viewing objects in poor light and can compensate for color changes in the lighting. Some types of image sensors have been developed to address these problems, but they can generate two to four times as much image data to process.

Image Processing

- Segmentation: image is broken down into objects of interest.
- Classification: objects are identified.
- Decision: objects are checked against rules and a decision is made.
- Output: the decision is output and acted on by the control system.

This is a simplified top level view of image processing. Some extra processing steps may be necessary depending on the type of problem that is being solved. In certain cases there may be multiple processing paths running in parallel where each process is optimized to solve a specific problem. For some problems a decision may not be necessary because the system only needs to record information for quality control.

Segmentation

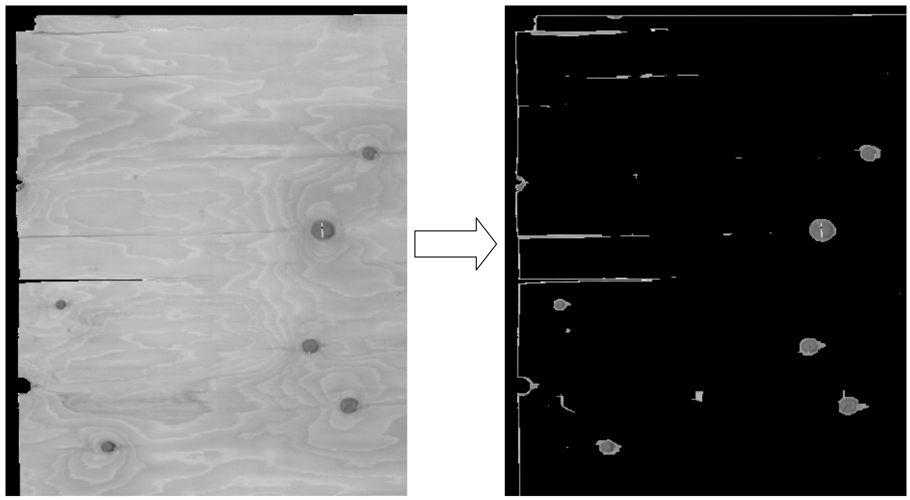

The basic idea of segmentation is to reduce the full image into areas of interest. This reduces the amount of image data that must be processed in detail. It is similar to the eye scanning process which also reduces the areas looked at.

This is only a simple example of segmentation. There are many ways to perform segmentation. The actual method used is problem dependent.

Classification

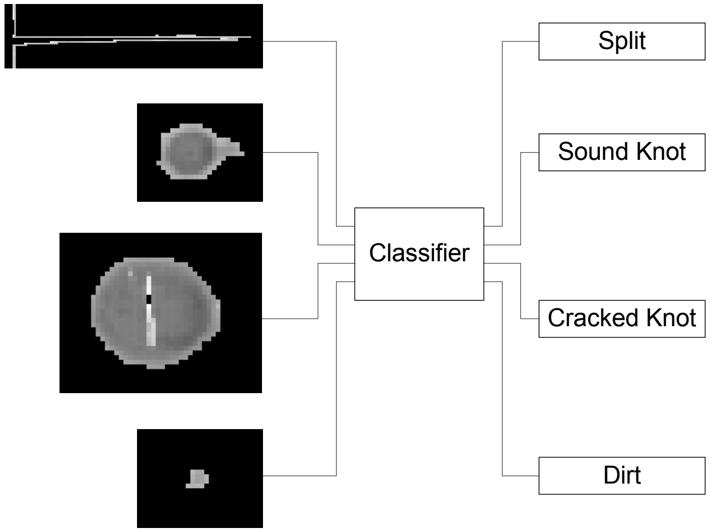

A classifier associates an object with a symbol. Humans constantly learn tasks of classification throughout their life cycle. The ability to assign symbols to objects is perhaps one of the things that make us human. We tend to learn this by example and with help from other humans. Superior machine vision classifiers attempt to model this process.

Decision

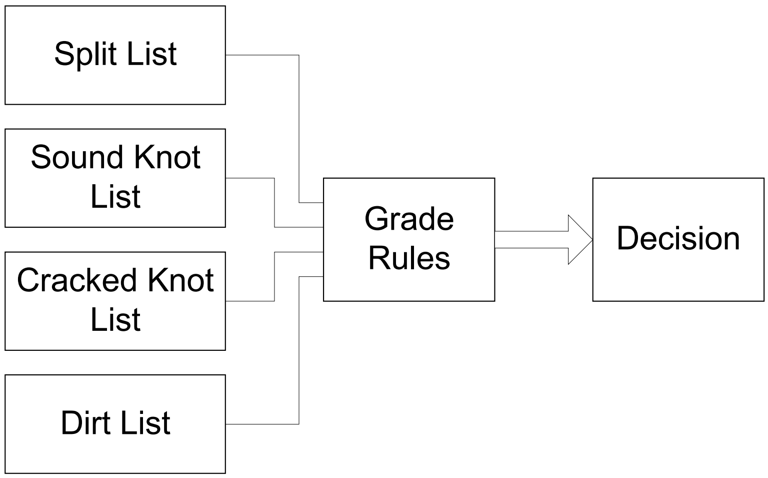

This is a simple example of a decision process. Not all the details are shown here.

The grade rules must be developed and entered into the machine by a human operator. Grade rules deal with specific types of defects and their maximum allowed size and number. This assumes a classifier is used to classify the defects into groups before the grade rules are applied.

In special cases, other factors such as material cost, product selling price, order backlog, inventory, manufacturing history and costs can be added into the decision process.

Machine Vision Process

- Camera sends only raw image pixels to processor.
- Line scan camera sees only one row of pixels at a time.
- Processors are primitive compared to eye and brain neural systems.
- Humans must still design and program vision system.

Classification Process

- Measurements (called *features*) are taken on the objects to be classified.
- In ideal case, different types of objects form clusters in measurement space.
- Classification is a statistical problem.
- No classifier is perfect, but they can be made nearly so through careful design.

Types of Classifiers

- Pattern matching: compares ideal template to objects.
- Rule based classifiers: hard or fuzzy pre-programmed rules separate classes.
- Learning classifiers: trained on samples of classes. Similar to rule classifier but more powerful.

These are the three basic groups of statistical classifiers. Most actual classifiers will fall into one of these groups. They all work by performing measurements on the object and then running these measurements though some mathematical or logical decision process. This is where statistics comes into the picture. Without proper statistical tools and measurements, classifier accuracy can never be known.

Pattern Matching

- Uses a correlation function between object template and image.
- Works well if object is uniform (example: computer chip on circuit board).
- Works poorly if object lighting, size or rotation is different than template.
- Mostly used on man made products (semiconductor chips and circuit boards).

Pattern matching assumes that the object will closely resemble the template. When the item is human or machine made, this may be true. In the case of natural objects there is usually too much variability, and the correlation function doesn't give a strong response.

Normally, pattern matching looks only at the object dimensions and grayscale information to make a decision. Higher level classifiers take more measurements into consideration.

Rule Classifier Example

- Feret is an image processing term.
- Feret X is width of object in pixels.
- Feret Y is height of object in pixels.

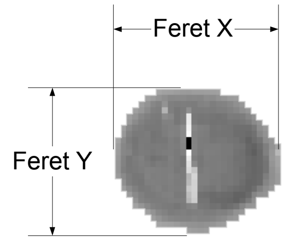

This simple example is used for both the rule classifier and the neural networks. The only reason the rule classifier is simple to understand is because we have limited the number of defect classes and inputs. Add the full range of defects and it won't work.

Rule Classifier Example

- Plot of Feret X and Feret Y for dirt (red), knot (blue) and stain (green) objects.
- Different classes form clusters of similar points.

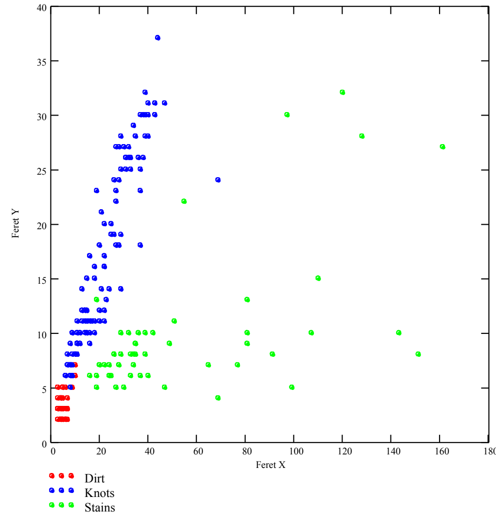

Rule Classifier Example

- Human expert draws decision lines or curves between classes.
- Lines are programmed into computer.
- Classifier operates using fixed set of rules.

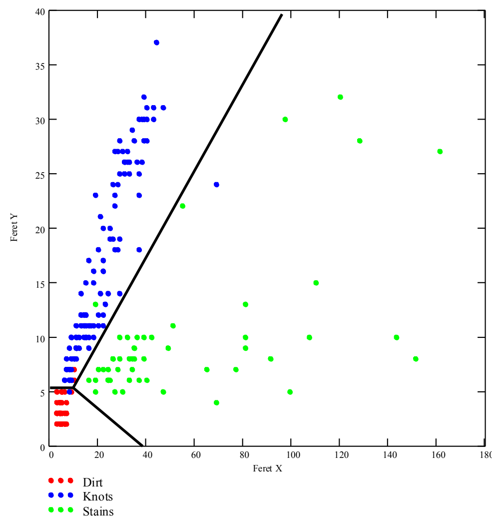

The upper left knot region can be compared to the neural network plots that will follow. It should also be noticed that one knot and one stain point are out of place and would not be classified properly by the rule classifier. These points that are out of place are called *outliners*.

Rule Classifier Drawbacks

- Requires tedious human effort to pick correct decision boundaries.
- Extension to more than two dimensions is difficult (requires multi-dimensional decision boundaries that are hard to visualize and implement in programs).
- Addition of a new class may require extensive redesign of decision tree.

Addition of more defect classes will cause this simple classifier to fail. The only solution to this problem is to use more measurements which will hopefully provide separation of the additional classes. Statistical tests will be needed to determine which measurements should be added.

It is not easy to represent these higher measurement dimensions on a plot and the decision surfaces must be selected out of the higher dimensional space. This is not very practical. A simple solution is to build decision trees of two dimensional classifiers in an attempt to capture the full set of dimensions. Neural networks can do this automatically and that is why they are used for these types of problems.

Artificial Neural Networks

- Neural networks are mathematical constructions that attempt to model some of the features of biological neural systems.
- Neural networks used as classifiers are a small subset of the full range of neural networks.
- The *Multi-Layer Perceptron* is the most common form of neural network classifier.

Biological nerve cells operate using electrical pulses and chemical signals. The frequency and strength of these pulses is processed by the nerve cells. Artificial neurons don’t operate on pulses but use similar concepts and perform similar operations when compared to real nerve cells.

Neural networks can be used for stock and commodity price prediction, data mining (trend analysis), credit risk and fraud detection, general purpose function fitting and digital filtering. There are types of networks that learn patterns presented to them and other types that can find their own patterns in the data.

Neural networks can be built out of analog computing elements, but they are easier to program and modify when simulated in a computer program.

The Artificial Neuron

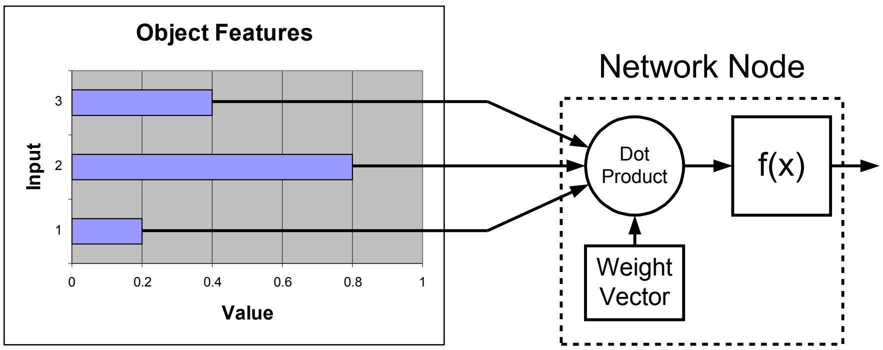

Artificial Neuron Function

- Inputs can be from either outside world measurements or other neurons.
- Dot product of input and weight vectors is calculated.
- Result is modified by non-linear function f(x) before output.
- Remove the non-linear function and you have a standard linear digital filter (FIR).

The study of neural networks and artificial neurons have many links to other fields. The fact that an artificial neuron can be easily converted into a standard digital filter makes them easier to understand when compared to the operation of a digital filter.

FIR stands for finite impulse response. It refers to the fact that the filter doesn’t use feedback and is stable under all input conditions. Filters that use feedback are called IIR (infinite impulse response) and can become unstable when fed certain sequences of inputs. Likewise, neural networks can use feedback and can also become unstable when used this way. The multi-layer perceptron is normally used in the feed-forward mode where there is no feedback.

Network of Artificial Neurons

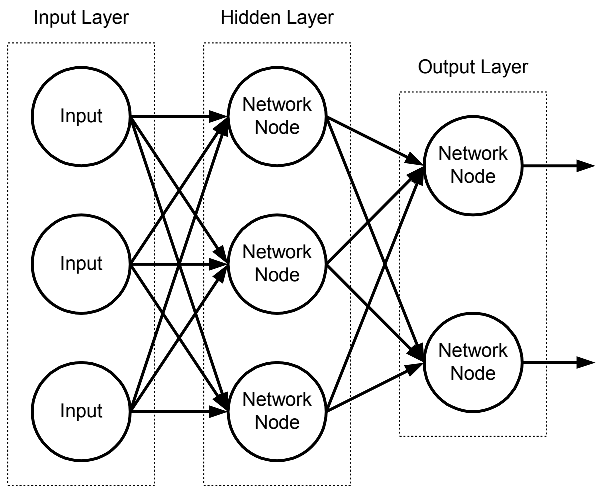

Networks can be built with any number of hidden layers. The number of hidden layers and the number of nodes in them determines how powerful the network is. For practical reasons, networks are usually limited to two hidden layers.

One example of how this network could work as a classifier:

(1) object measurements are entered into the inputs.

(2) all network nodes are calculated in a forward moving process.

(3) the output with the highest value corresponds to the class of the object.

Network of Artificial Neurons

- Connected networks of neurons form a neural network.
- Mathematically, it performs a non-linear mapping from the n-dimensional input space to the output space.
- The form of this mapping function can be trained into the network using the back-propagation algorithm.

Artificial neurons are only useful for classification problems when connected together into a network. The mapping function of a single neuron is too simple for all types of classification problems.

Neural networks can be studied for what they are: large mathematical constructions. Linear algebra, non-linear optimization and calculus are the main fields used.

The back-propagation algorithm made neural networks practical. Previously, there was no easy way to design a network for a given problem. Many other optimization methods can be used to train neural networks, but back-propagation is the simplest and easiest to use.

Network Mapping Functions

- A single layer can perform simple *AND* and *OR* Boolean logic functions.
- Two layers can perform the more complex *XOR* Boolean logic function.
- Three or more layers can perform even more complex mappings.

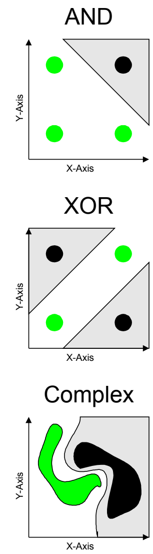

A single neuron can only separate objects into two classes using a straight line decision boundary. The values of the weight vector determine the placement of the decision boundary. When another layer of neurons is added, it is possible to separate objects into multiple areas and use a curved decision boundary. Add a third layer and arbitrary decision shapes can be computed. The complexity of the decision shape is only limited by the number of nodes that can added. Larger networks, with many layers and nodes, are harder to train.

Network of Artificial Neurons

- Neural networks are similar to statistical regression methods.
- Replace the non-linear functions with linear functions and the network collapses into a linear regression problem.
- Now for some examples using the previous veneer data...

The back-propagation training method is based on least squares regression. While the goal of both is minimization of the least squares error function, neural networks are not linear systems and can't be solved using the same methods used for linear regression. They must be solved using optimization methods, of which back-propagation is only one example.

1 Hidden Layer, Normal Training Knot Region Selected

- Deep red is knot area. Blue is dirt & stain.
- Doesn't fit all points.
- Odd structure in region of knot & stain mixing.

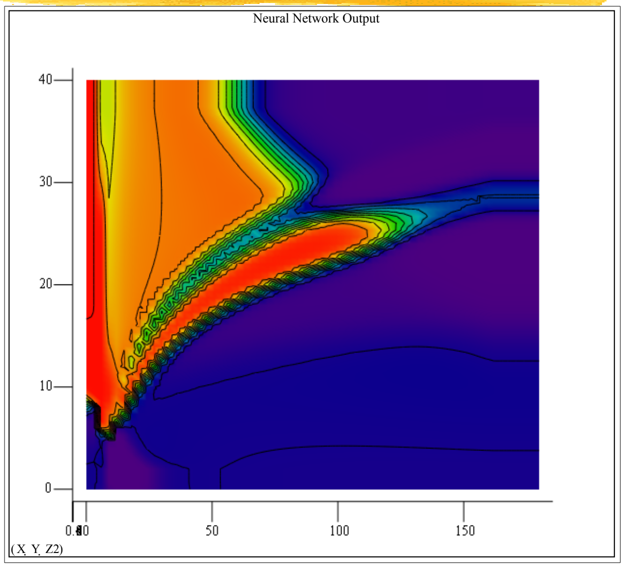

This is a one hidden layer network trained with normal methods.

The graph is 3-dimensional with the Z axis pointing up out of the slide. It is looking directly down on the response of the knot output neuron when plotted over feret x and feret y. Different levels of the function are represented by rainbow colors from red to blue. Contour lines show steps of 0.1 from 0 (bottom blue) to 1 (top red). The small ripples are caused by the increments in the plot grid, the actual neuron output response is smoother than shown.

This red area can be compared to the rule classifier knot region which was a simple trapezoid. There is some mixing of class points in this example. As the network tries to fit these outliner points it will generate odd shapes and groves.

2 Hidden Layers, Normal Training Knot Region Selected

- Still doesn't fit all points.
- Single knot point selected by extended region.
- Small error spike near dirt at bottom.

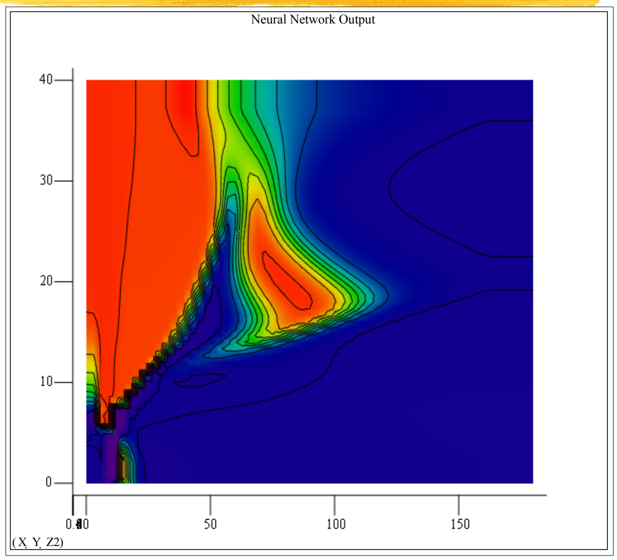

This network has two hidden layers and can fit more complex shapes. It has been able to fit a single knot point that is in the stain region.

1 Hidden Layer, Special Training Knot Region Selected

- Fits all points.
- Hole near bottom fits single stain point.
- Extra holes provide poor fit in knot region.

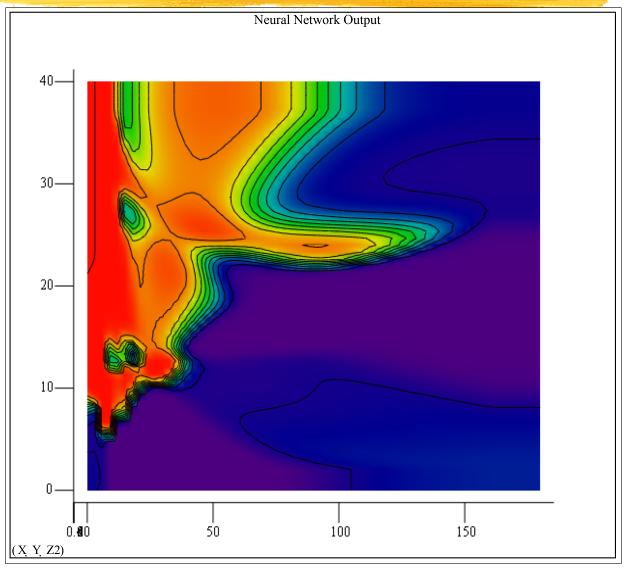

A special version of back-propagation was used to train a one hidden layer network to fit all the points. As a result of making the network fit all the points, some extra holes were introduced into the knot region.

2 Hidden Layers, Special Training Knot Region Selected

- Fits all points.
- No extra holes.
- Large stain regions selected.

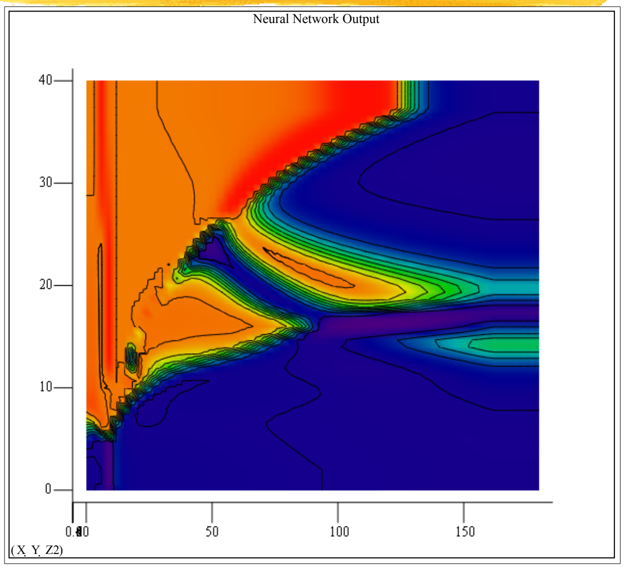

Here, the special training was used to make a two hidden layer network fit all the points. While it has fit all of the knot region correctly, it has extended large selection areas into the stain region.

Training Conclusions

- Outliner points (examples outside their normal range) can cause improper selections of class regions.
- Regions not nailed down by examples can end up in any class.
- Statistical tests and restrictions must be applied to the training data to prevent errors in classification.

The whole point of these examples was to show that using raw training data can produce a correct classifier but give a poor statistical result. The network can be made to fit all the training points correctly but it can give poor results when used to classify new data points.

1 Hidden Layer, Normal Training Selected Training Data Points

- No odd regions extending into other classes.
- Still needs more data points added for best accuracy.

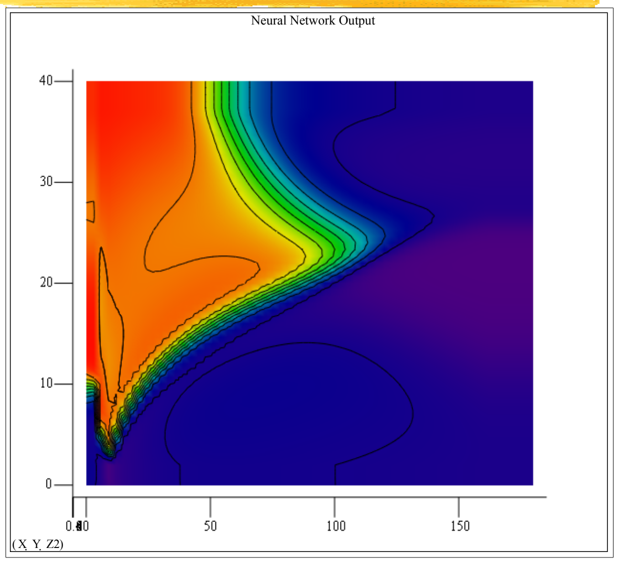

In this case we have removed the inconsistent training points from the data. This has eliminated the strange behavior of the network and produced a better fit to the data. It still is selecting a little too much of the stain region, suggesting that we need to add more stain examples.

Conclusions

- Neural networks make very powerful statistical classifiers when used properly.
- Neural networks can automatically solve complex classification problems.
- Real problems involve more defect classes, input dimensions and complex network structures than shown by this simple example.