Bruno ANGELES
Slide presentation II
8 October 2009

Artificial Neural Networks

Introduction

Artificial neural networks (ANNs) are a structure used in computer science and electrical engineering, based on biological neurons. They have applications in video games, medicine, robotics, marketing, and fraud detection, among other fields.

Biological Neurons

A biological neuron's nucleus sums the inputs that come from the synapses through the axon, and the resulting signal is output through the dendrites. A neuron's output is either on or off, and the logic that dictates this triggering is called the activation function. When a certain excitation is applied repeatedly to a neuron by another one, the strength of this connection is strengthened and it becomes easier for the same excitation to occur. Plasticity happens when a neuron is not stimulated for a long time by another neuron: the effectiveness of the connection decreases. Human neurons' rate of transmission is of the order of 100 Hz, due to the chemical to electric transduction at the synapses. Despite this low frequency rate compared to contemporary processors, systems using neural networks (e.g. the human brain) can outperform computers at many tasks. Artificial neural networks are an attempt at modeling the structure of such systems in order to solve problems that may not be solvable otherwise.

To the Digital World

Digitally, neurons are modeled by a set of weighted inputs, an adder that performs the weighted sum of inputs, and an activation function. The unit function that serves as the model of biological neurons' activation function is replaced by an S-shaped curve, often the sigmoid function $y(t) = \frac{1}{1 - e^{-\beta t}}$, where β > 0, because it is conveniently derivable (and its first derivative can be evaluated), a useful feature in the backpropagation algorithm used in the supervised training of ANNs.

The most common form of ANNs is feed-forward networks. In feed-forward networks, all arrows of signal propagation go towards the output. A feedback network is a network in which at least a connection goes from a neuron's outputs to an earlier neuron's inputs in the sequence from inputs to outputs. ANNs need not have only one output. To certain users, a neural network might operate as a black box system, in which only the number of inputs and outputs is known. This can be seen as both an advantage and a disadvantage.

Applications

ANNs are often implemented when data is available for a certain problem but no theory is, when the input data has no obvious pattern that can be used for modeling, or when robustness to noise is of particular importance. Applications of this computing method exist in object recognition, medical diagnosis, obstacle avoidance, environment exploration, sales forecasting, marketing, and fraud identification. In Music Technology and Speech Technology, they have been used to recognize human speakers, perform text-to-speech conversion, and transcribe polyphonic music.

Training an ANN is done in one of three ways: unsupervised learning, supervised learning, and reinforcement learning. It is meant to update the weights of the system in order to improve the behaviour of the ANN.

Training Methods

Unsupervised training works by minimizing a given cost function, dependent on the inputs and outputs of the system. The cost of a compression system could be the ratio of compression, and the cost of a filter could be its signal-to-noise ratio. No known data set can minimize the cost function, so the system is fed a set of data and finds data that minimizes the cost function.

Supervised training works by using a teacher to check the error of the ANN and minimize it, knowing the expected output. The most popular method of supervised training is backpropagation. This concept was introduced by Paul Werbos in 1974, but wasn't quite used until about a decade later, with a resurgence of interest in ANNs. Two issues with supervised training are being stuck at a local minimum and overfitting. The former is a situation in which a solution is found that is not the global solution, although it may seem like it, while the latter means that the ANN has been trained to respond to a certain input set so well that it will lose its capacity to generalize. Local minima can be avoided by adding momentum to the weight update of the nodes, and overfitting can be avoided by either reducing the number of neurons in the system, adding noise to the input (e.g. alter the input fonts to an optical character recognition system), or regularly verifying the ANN's behaviour by using validation data in parallel to the regular training data.

Backpropagation works on feed-forward networks by evaluating the error δ of the system at the output of the ANN. The final error is propagated backwards through each layer of the ANN, giving us an evaluated error at each node. Each weight of the network is updated by a product of a learning speed, the final error δ, the current node output and the first derivative of the activation function, evaluated at the node's current input. In order to implement momentum, the previous weight update must also be part of the current weight increment.

The third method of training for an ANN is reinforcement training. In this method, a long-term reward objective is pursued. At each frame, the state of the ANN is analyzed and the action that maximises the final output is chosen.

These training methods are intended to adjust the behaviour of a system that may have been initialized with purely random weights.

Conclusion

On the one hand, artificial neural networks have the advantage of being easy to implement, and flexible in certain cases. The fact that they can learn could also be appealing to certain applications, such as toys that recognize their handlers. On the other hand, there are also certain pitfalls that must be avoided, and the black box model might not satisfy all users. It should also be remembered that artificial neural networks can always be combined to other computing methods as part of a global system.

References

Buckland, Mat. 2002. *AI Techniques For Game Programming*: Premier Press.

Karaali, O., G. Corrigan, N. Massey, C. Miller, O. Schnurr, and A. Mackie. 1998. A high quality text-to-speech system composed of multiple neural networks. *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing* 2:1237-40.

Marmanis, H., and D. Babenko. 2009. *Algorithms of the Intelligent Web*: Manning Publications.

Marolt, M. 2001. Transcription of polyphonic piano music with neural networks. *Proceedings of the 10th Mediterranean Electrotechnical Conference, 2000. MELECON 2000.* 2:512-5.

Murray, J. C., H. R. Erwin, and S. Wermter. 2009. Robotic sound-source localisation architecture using cross-correlation and recurrent neural networks. *Neural Networks* 22 (2):173-89.

Rho, S., B. Han, E. Hwang, and M. Kim. 2007. MUSEMBLE: A Music Retrieval System Based on Learning Environment. *Proceedings of the 2007 IEEE International Conference on Multimedia and Expo*:1463-6.