Neural Networks for Intelligent Vehicles

Dean A. Pomerleau
Carnegie Mellon University, School of Computer Science
5000 Forbes Ave., Pittsburgh, PA 15213-3890
Phone: (1)412-268-3210, Fax: (1)412-621-1970, Internet: pomerleau@cmu.edu

Abstract

This paper is a survey of current research in applying artificial neural networks to the domain of intelligent vehicles. It describes work in three areas: video-based traffic monitoring, monitoring and control of onboard systems, and vision-based lateral control. In each of these domains, successful preliminary systems demonstrate that artificial neural networks have the potential to make significant improvements in the state-of-the-art. Because of the simplicity and uniformity of the neural network architectures and algorithms employed in these systems, they each have the potential to be implemented efficiently in hardware, which could eventually make them commercially viable.

1. Introduction

Artificial neural networks are a powerful technique for solving non-linear mapping and classification tasks. They have proven capable of solving problems ranging from speech and handwritten character recognition [29] [26] to medical diagnosis [2]. Recently neural networks have even begun to appear in consumer products such as washing machines and vacuum cleaners [17].

Two attributes of artificial neural networks make them useful for a wide variety of problems, including intelligent vehicles. The first beneficial attribute is their ability to adapt to new tasks and circumstances with relatively little effort on the part of the user or developer. In their most common form, namely multi-layer perceptrons, neural networks learn to map input patterns to particular output patterns or classes based on many examples of the desired mapping. For instance, the input pattern in character recognition is typically a bitmap image of a character, and the output is the identity of the character. The system developer need not entirely specify the relevant features or processing required for the task, but instead needs only provide the network with numerous examples of the mapping to be performed (e.g. bitmaps paired with their identities). Neural network training algorithms, such as back-propagation [25] or radial basis function learning [19] can automatically determines from these examples the relevant input attributes and how to process them to perform the task. This ability to learn complex tasks through observation makes neural networks well suited to domains like intelligent vehicles, where the problems are often ill understood and constantly changing.

The second important attributed of artificial neural networks is their computational simplicity. In general, artificial neural networks are composed of many very simple computing elements, called units, which interact with each other through weighted connections. The simplicity and uniformity of this computing paradigm makes efficient implementations possible on both serial and parallel computers [13] [33] [31]. In addition, the regular nature of the neural network processing makes hardware implementations particularly attractive. The ability to implement a trained neural network on a single chip both reduces cost and increases reliability, two crucial factors in determining commercial viability. Nowhere is the need for inexpensive, reliable systems more acute than in the domain of intelligent vehicles, where the potential market includes the millions of cars sold each year, and where the cost of failure can be catastrophic.

This paper discusses three applications of artificial neural networks to the domain of intelligent vehicles: video-based traffic monitoring, monitoring and control of onboard systems, and vision-based lateral control. None of these systems is yet commercially available, but each has the potential to greatly improve the safety and efficiency of road travel.

2. Traffic Monitoring

The task in traffic monitoring is to detect traffic level and flow rates at key points, such as traffic lights and tunnels. The most common approach to traffic monitoring is install mechanical or electrical devices placed on top or embedded in the roadway. There are three disadvantages to such systems. First, they are expensive to install, particularly if they must be embedded in the pavement like the electrical inductance "loop" detectors typically used at traffic lights. Second, they are prone to mechanical failure due to pavement cracking. Third, they often can only "see" a single lane of traffic, or if they can detect traffic in multiple lanes, can not discriminate traffic on one lane versus another.

More recently, people have begun to investigate the use of machine vision techniques for traffic monitoring [5] [28]. These systems have the advantages of being easy to install, and being able to provide much more specific information about traffic conditions, including queue length and lane distribution. These systems generally employ simple image processing like image thresholding, image differencing, and template matching. As a result, they suffer from several problems, including sensitivity to lighting conditions and camera perspective. While more sophisticated image processing system which use optical flow and appearance models for different vehicle types are under development, they generally suffer from brittleness under changing

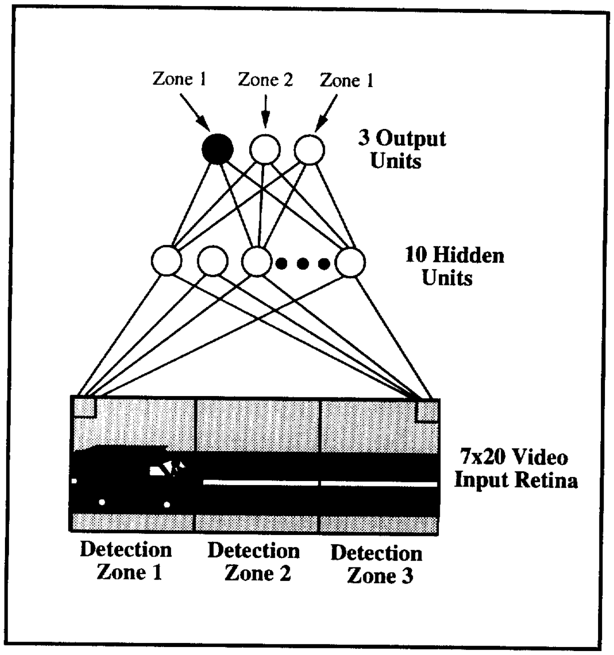

Figure 1: Network for Vehicle Detection.

conditions.

An alternative which may prove more robust is to use neural network-based image processing. Just such a system is being developed by Bullock, Garrett and Hendrickson [6]. The idea is to train a network to monitor a scene and indicate when the image contains a vehicle and where the vehicle is located. The network receives a reduced resolution video image of the scene as input, and has a number of output units each trained to respond to the presence of a vehicle in a particular region of the image (See Figure 1). The particular architecture used in this experiment had a 7x20 pixel input retina, 10 hidden units, and 3 output units corresponding to three "detection zones".

The network was trained on 233 images digitized from a two hour videotape of a public street. The images contained large amounts of visual noise including heavy shadows, bicycles, pedestrians and joggers, as well as many types of vehicles including cars and trucks of different makes, models and colors. The network was trained using a variation of the back-propagation learning algorithm called quickprop [10].

The trained network was tested on a disjoint test set of 83 images digitized from the same videotape. The network's performance was measured using the following criteria. The network was considered to have made a mistake if:

- a vehicle is present in a zone, but the corresponding output unit has an activation less than 0.6
- a vehicle is not present in a zone, but the corresponding output is greater than 0.4
- one of the outputs is between 0.4 and 0.6, a range considered indeterminant.

Using this criterion, the network had an error rate of 12However all 9 of the mistakes occurred when a vehicle was transitioning from one zone to another. All 9 error were eliminated when the temporal activation pattern of the output units was examined.

These results are very preliminary, in that the system does not take into account multiple vehicles in the scene at once, and do not address the problem of vehicle identification (cars vs. trucks). However they do demonstrate that neural networks are capable of effective traffic monitoring, even in the presence of changing lighting conditions.

3. Vehicle Control Systems

A second successful application of artificial neural networks to the domain of intelligent vehicles has been in the area of monitoring, diagnosis and control of vehicle control systems. Increasing demands for performance, safety and fuel economy and lead to increasingly complex and sophisticated vehicle control systems in today's automobiles. A few of these advances include electronic ignition, anti-lock braking and active suspension systems. To achieve the full benefit of these advances requires sophisticated monitoring and control. What makes these applications particularly challenging is that they require very rapid responses, often based on quite limited amounts of data. Artificial neural networks have demonstrated an ability to perform well within these severe constraints.

3.1. Ignition Timing Estimation

A prime example of the effect use of neural networks for vehicle control system monitoring is the work of Willson, Whitham and Anderson [30]. They have focused on the task of neural network based engine control using information such as cylinder pressure, manifold pressure, engine temperature, engine speed and engine emissions. Their preliminary work has focused on on-line estimation of ignition timing, a crucial factor contributing to fuel economy and vehicle performance. With the proper equipment, ignition timing is straightforward to measure. Unfortunately such equipment is difficult to integrate into a moving vehicle. Recent development of high speed piezoelectric and fiber optic pressure sensors have made it possible to indirectly measure ignition timing by recording cylinder pressure at various crank angles. However the relationship between the time course of cylinder pressure and ignition timing is highly non-linear due to its dependence on engine speed, air/fuel ratio and throttle setting.

Willson, Whitham and Anderson employed a single hidden layer network similar in many respects to the network used for traffic monitoring. However in their network, they employed 36 input units represented representing cylinder pressure at 36 different crank angles between $-40°$ and $+40°$ of top dead center. The best network architecture they found for the task had 16 hidden units and a single output unit representing the spark timing, which ranged from $5°$ to $35°$ before top dead center.

The network was trained on cylinder pressure data collected from a 350 $in.^3$ Chevrolet engine running at 4000 RPMs with various spark timings. After training, the network was able to estimate the ignition timing on a disjoint test set collected from the same engine to within an average of 0.66°. While quite encouraging, the network's preformance under different conditions (e.g. different engine RPM or air/fuel ratio) remains to be tested.

3.2. Anti-lock Brake System Control

Another area of vehicle controls where early experiments suggest neural networks can be effectively applied is in the control of anti-lock braking systems [9]. In this task, the goal is to stop the vehicle as quickly as possible. As anyone who has driven on ice or snow realizes, the optimum braking strategy is not to slam on the brakes. This response will cause the wheels to "lock-up" and slide with much less friction than will occur when moderate braking is applied. The task is complicated by the fact that braking performance is heavily dependent on road conditions and true vehicle velocity, factors which are difficult to measure accurately.

The Davis et al. anti-lock braking system employs a two stage training system typical of neural network control systems. In the first stage, a "plant identification" network is trained to predict the behavior of the physical system under a variety of conditions. They trained a recurrent network to predict the wheel and vehicle velocity in the next time step based on the current wheel and vehicle velocity, and the current braking command. The training algorithm used was the dynamic decoupled extended Kalman filter (DDEKF) algorithm described in [11].

Once the identification network was trained to predict the vehicle's behavior, it was used to train a second recurrent network to produce the braking command for optimal stopping. The controller network learned from the vehicle model just how much braking force can be applied before the wheels lock under a variety of road conditions and vehicle speeds. After training in simulation using the identification network as the vehicle model, the controller network was trained on-line using a real vehicle driven on a test track. The controller network was also trained using the DDEKF algorithm.

Results showed that the network was able to learn to keep the vehicle's brakes from locking and the wheel slippage quite close to the theoretical optimal under a variety of road conditions (coefficients of friction between 0.2 to 1.0). While this performance is encouraging, Davis et al. did not report any comparison between the stopping distance achieved by the neural network controller and existing anti-lock braking systems.

Similar promising results have been achieved in active suspension control and detection of engine misfires [15]. In the later domain, a feedforward multi-layer perceptron was able to detect all but four of 150 misfires in a test set of 7600 engine measurements recorded from a vehicle in a wide variety of conditions (potholed roads, rapid acceleration and braking). It exhibited only 13 false alarms on the 7450 normal engine readings. This was in fact the first successful automated misfire detection system operating on a vehicle.

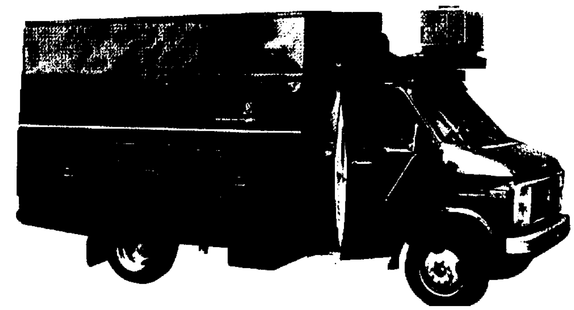

Figure 2: The CMU Navlab autonomous navigation testbed vehicle.

4. Vision-based Lateral Control

Perhaps the most widely studied application of neural networks to intelligent vehicles has been in vision-based lateral control. The goal in this task is to steer the vehicle based on input from an onboard video camera.

There are a number of neural network-based vision systems for autonomous vehicle control. Some rely primarily on color information [32] [16] to determine where the road is in the scene, while others use texture variations [7] in the image of the scene ahead of the vehicle to make steering decisions. The neural network architectures and algorithms employed by these systems also varies considerably. Many rely on standard feedforward networks and the back-propagation algorithm for training [32] [16], while others use recurrent networks [7] or radial basis functions [1].

The first and most successful neural network-based autonomous driving system is the ALVINN system developed by Pomerleau [24]. The ALVINN system is designed to drive the CMU Navlab autonomous navigation testbed shown in Figure 2. The vehicle is equipped with a video camera, and motors on the steering wheel, brake and accelerator pedal, enabling computer control of the vehicle's trajectory.

4.1. Neural Network Model

The connectionist model for autonomous road following used in the ALVINN system [23] is the feedforward multi-layer perceptron shown in Figure 3. The input layer consists of a single 30x32 unit "retina" onto which a video image is projected. Each of the 960 input units is fully connected to the four unit hidden layer, which is in turn fully connected to the output layer. The 30 unit output layer is a linear representation of the currently appropriate steering direction. The centermost output unit represents the "travel straight ahead" condition, while units to the left and right of center represent successively sharper left and right turns.

To drive the Navlab, an image from the video camera is reduced to 30x32 pixels and projected onto the input layer. After propagating activation through the network, the output layer's activation profile is translated into a vehicle steering command.

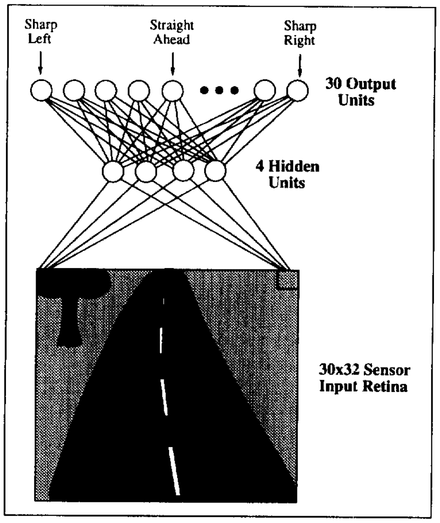

Figure 3: Architecture of the network designed for autonomous driving.

4.2. Training "On-the-Fly"

The most interesting and novel aspect of the ALVINN system is the method used to train it. In this technique, called training "on-the-fly" the network is taught to imitate the driving reactions of a person. As a person drives, the network is trained with back-propagation using the latest video image as input and the person's steering direction as the desired output.

To facilitate generalization to new situations, variety is added to the training set by shifting and rotation the original camera image in software to make it appear that the vehicle is situated differently relative to the road ahead. The correct steering direction for each of these transformed images is created by altering the person's steering direction for the original image to account for the altered vehicle placement. So for instance, if the person was steering straight ahead, and the image was transformed to make it appear the vehicle is off to the right side of the road, the correct steering direction for this new image would be to steer towards the left in order to bring the vehicle back to the road center. Adding these transformed patterns to the training set teachs the network to recover from driving mistakes, without requiring the human trainer to explicitly stray from the road center and then return.

4.3. ALVINN Driving Performance

Running on two Sun Sparcstations onboard the Navlab, training on-the-fly requires about two minutes during which a person drives over about a 1/4 to 1/2 mile stretch of training road. During this training phase, the network typically is presented with approximately 50 real images, each of which is transformed 15 times to create a training set of 750 images.

Once it has learned, the network can accurately traverse the length of the road used for training, and also generalize to drive along parts of the road not encountered during training under

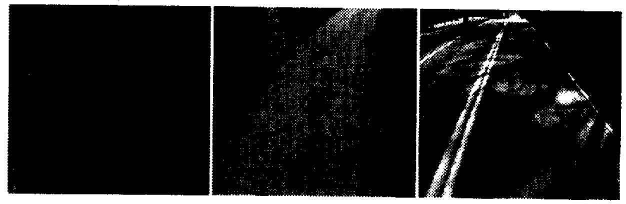

Figure 4: Video images taken on three of the roads ALVINN has been trained to handle.

a variety of weather and lighting conditions. In addition, since determining the steering direction from the input image merely involves a forward sweep through the network, the system is able to process 20 images per second, and drive at up to 55 mph. This is over five times as fast as any non-connectionist system as driven using comparable hardware [14] [8].

The flexibility provided by the neural network has allowed ALVINN to learn to driving in a wide variety of situations. Individual networks have been trained to drive on single-lane dirt and paved roads, two-lane suburban and city streets, and multi-lane divided highways. Images taken from three of these domains are shown in Figure 4. On the highway, ALVINN has driving for for up to 21 miles without human intervention. Measurements of the ALVINN systems driving accuracy show it is able to keep the vehicle within 6.9cm of the center of its lane on average [22]. This compares favorably with human driving performance, which has been measured to average 5.7cm from the center of the lane on average [3].

The eventual goal of the ALVINN project is to create an advanced cruise control system which controls both longitudinal and lateral motion of the vehicle. The system will be trained by watching the person drive for several minutes. After training, the person will push a button and ALVINN would take over all vehicle controls. However this fully autonomous system is still a long way off. In the nearer term, the ALVINN system will be used as a lane excursion warning device, sounding an alarm when the driver starts to swerve out of his lane. In fact, a single chip hardware implementation of the ALVINN system, integrating both the CCD array for image capture and the neural network for image processing, is currently under development at Carnegie Mellon University. This hardware implementation should provide an inexpensive and reliable platform on which to develop a practical lane excursion warning device.

5. Conclusion

Artificial neural networks have the potential to make large contributes to the field of intelligent vehicles. Their ability to learn complex non-linear mappings allow neural networks to perform many difficult intelligent vehicle tasks, including traffic monitoring, mechanical system monitoring and control, and vision-based lateral motion monitoring and control.

While such applications of neural networks are still in the experimental stages, their preliminary success, when coupled with the fact that their simplicity allows them to be implemented cheaply and reliably, should result in commercially viable intelligent vehicle products. Such commercial neural networks are already available in certain consumer products including models of washing machines, vacuum cleaners and air conditioners sold by Panasonic [17]. A single chip implementation of the ALVINN neural network vision system is currently under development for use as a lane excursion warning system.

However there are several factors currently limiting the widespread application of neural networks to real world problems, including those of intelligent vehicles. The first is a general lack of rigorous performance comparisons with alternative methods. While there has been some limited experiments comparing the driving accuracy of the ALVINN system with that of people, in general intelligent vehicle applications of neural networks have only reported qualitative results. Part of this lack of rigor stems from the fact that it is often difficult to find good metrics for performance in domains like traffic monitoring, and once they are found, it is difficult to measure a system's performance. For example, measuring the accuracy of the ALVINN driving system involved painstaking experiments in which water was dripped from the center of the vehicle and then the position of the water drops were measured relative to the road center. While such quantitative measurements are difficult to perform, they are absolutely necessary if neural network solutions to problems are to be accepted.

The need for rigorous performance characterization is increased by the inherent difficulty in analyzing artificial neural networks. Neural networks are computationally a very powerful technique, in fact they are able to represent arbitrarily complex real-valued mappings from inputs to outputs [4]. But as a result of this power, the internal representations they develop are often quite hard to interpret, and the processing they perform can be quite difficult to understand. A number of techniques are under develop which can provide greater insight into the processing and representations of neural networks, including cluster analysis [12], sensitivity analysis [21], explanation based neural networks [18], and reliability estimation techniques [20]. When perfected and widely employed, these techniques should mitigate many of the concerns over the "black-box" nature of neural networks and allow them to be applied to critical real world tasks like intelligent vehicles.

6. Acknowledgements

The principle support for the Navlab has come from DARPA, under contracts DACA76-85-C-0019, DACA76-85-C-0003 and DACA76-85-C-0002. This research was also funded in part by a grant from Fujitsu Corporation.

References

[1] Aste, M., and Caprile, B. (1992) Learning autonomous navigation abilities using radial basis function networks. In *Proceedings of the 1992 Intelligent Vehicles Symposium*, I. Masaki (ed.), pp 241-246.

[2] Baxt W. (1993) The Application of the artificial neural network to clinical decision making. *Advances in Neural Information Processing Systems 5*, Giles, C.L., Hanson, S.J., and Cowan, J.D. (ed.), Morgan Kaufmann.

[3] Blaauw, G.J. (1982) Driving experience and task demands in simulator and instrumented car: A validation study, *Human Factors, Vol. 24*, pp. 473-486.

[4] Baldi, P. (1991) Computing with arrays of bell-shaped and sigmoid functions. *Advances in Neural Information Processing Systems 3*, R.P. Lippmann, J.E. Moody, and D.S. Touretzky (ed.), Morgan Kaufmann, pp. 735-742.

[5] Blosseville, J., Krafft, C., Lenoir, F., Motyka, V., and Beucher, S. (1989) TITAN: A traffic measuring system using image processing techniques. *Second Int. Conf. on Road Traffic Monitoring*, pp. 84-88.

[6] Bullock, D., Garrett, J. and Hendrickson C. (1991) A prototype neural network for vehicle detection. *Artificial Neural Networks in Engineering*.

[7] Catala, A., Grau, A., Morcego, B., Fuertes, J. (1992) A neural network texture segmentation system for open road vehicle guidance. In *Proceedings of the 1992 Intelligent Vehicles Symposium*, I. Masaki (ed.), pp 247-252.

[8] Crisman, J.D. and Thorpe, C. (1990) Color Vision for Road Following. *Vision and Navigation: The CMU Navlab* C. Thorpe (Ed.), Kluwer Academic Publishers, Boston.

[9] Davis, L., Puskorius, G., Yuan, F., and Feldkamp L. (1992) Neural network modeling and control of an anti-lock brake system. In *Proceedings of the 1992 Intelligent Vehicles Symposium*, I. Masaki (ed.), pp 179-184.

[10] Fahlman, S.E. (1988) Faster-learning variations on back-propagation: an empirical study. *Proceedings of the Connectionist Summer School*, Morgan Kaufmann.

[11] Feldkamp, L., Puskorius G., Davis, L., and Yuan, F. (1991) Decoupled Kalman training of neural and fuzzy controllers for automotive systems. *Proc. of the Fuzzy and Neural Systems and Vehicle Applications Conference*, Tokyo, Japan.

[12] Gorman, R.P. and Sejnowski, T.J. (1988) Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks Vol. 1*, pp. 75-89.

[13] Gusciora, G.L., Pomerleau, D.A., Touretzky, D.S., Kung, H.T. (1990) Back-propagation on Warp. *Artificial Neural Networks: Applications and Implementations*, Ben Wah, Manoel F. Tenorio, Pankaj Mehra and Jose A.B. Fortes (Eds.) IEEE Computer Society Press.

[14] Kluge, K. and Thorpe, C. (1990) Explicit models for robot road following. *Vision and Navigation: The CMU Navlab* C. Thorpe (Ed.), Kluwer Academic Publishers, Boston.

[15] Marko, K.A. (1991) Neural network application to diagnostics and control of vehicle control systems. In *Advances in Neural Information Processing Systems 3*, R.P. Lippmann,

J.E. Moody, and D.S. Touretzky (ed.), Morgan Kaufmann, pp. 537-543.

[16] Marra, M., Dunlay, T.R., Mathis, D. (1988) Terrain classification using texture for the ALV. Martin Marietta Information and Communications Systems technical report 1007-10.

[17] Maruno S. (1993) Smart consumer products using neural networks in Japan. Invited talk *Neural Networks for Computing*, Snowbird, Utah.

[18] T.M. Mitchell and S.B. Thrun (1993) Explanation-based neural network learning for robot control. To appear in *Advances in Neural Information Processing Systems 5*, Giles, C.L., Hanson, S.J., and Cowan, J.D. (eds.) Morgan Kaufmann.

[19] Poggio, T., and Girosi, F. (1990) Regularization algorithms for learning that are equivalent to multilayer networks. *Science, Vol. 247* pp. 987-982.

[20] Pomerleau, D.A. (1993) Input Reconstruction Reliability Estimation. *Advances in Neural Information Processing Systems 5*, Giles, C.L., Hanson, S.J., and Cowan, J.D. (eds.) Morgan Kaufmann.

[21] Pomerleau, D.A. and Touretzky, D.S. (1993) Understanding Neural Network Internal Representations through Hidden Unit Sensitivity Analysis. *Proceedings of the International Conference on Intelligent Autonomous Systems* , C.E. Thorpe (ed.), IOS Publishers, Amsterdam.

[22] Pomerleau, D.A. (1992) Neural network perception for mobile robot guidance. PhD. Dissertation. Carnegie Mellon technical report CMU-CS-92-115.

[23] Pomerleau, D.A. (1991) Efficient Training of Artificial Neural Networks for Autonomous Navigation. In *Neural Computation 3:1*.

[24] Pomerleau, D.A. (1989) ALVINN: An Autonomous Land Vehicle In a Neural Network. *Advances in Neural Information Processing Systems 1*, D.S. Touretzky (ed.), Morgan Kaufmann.

[25] Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986) Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. I: Foundations* pp. 318-362, Bradford Books/MIT Press, Cambridge, MA.

[26] Schenkel, M., Weissman, H., Guyon, I., Nohl, C., Henderson, D., Boser, B., and Jackel, L (1993) *Advances in Neural Information Processing Systems 5*, Giles, C.L., Hanson, S.J., and Cowan, J.D. (ed.), Morgan Kaufmann.

[27] Seitzma, J., and Dow, R. (1991) Creating Artificial Neural Networks that Generalize. *Neural Networks, Vol. 4* pp. 67-79, Pergamon Press.

[28] Shimizu, K., and Shigehara, N. (1989) Image processing system using cameras for vehicle surveillance. *Second Int. Conf. on Road Traffic Monitoring*, pp. 61-65.

[29] Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., and Lang K. (1987) Phoneme recognition using time-delay neural networks. ATR Technical Report TR-I-0006.

[30] Willson B., Whitham, J., and Anderson C. (1992) Estimating ignition timing from engine cylinder pressure with neural networks. In *Proceedings of the 1992 Intelligent Vehicles Symposium*, I. Masaki (ed.), pp 108-113.

[31] Witbrock, M. and Zagha, M. (1989) An implementation of back-propagation learning on the GF-11, a large SIMD parallel computer. Carnegie Mellon University tech report CMU-CS-89-208.

[32] Wright, W.A. (1989) Contextual road finding with a neural network. British Aerospace Advanced Information Processing Department technical report.

[33] Zhang, X., Mckenna, M., Misirov, J., and Waltz, D. (1990) An efficient implementation of the back-propagation algorithm on the connection machine CM-2. *Advances in Neural Information Processing Systems 2*, D. Touretzky (ed.), Morgan Kaufmann.