Ειδκά Συστήματα Ελέγχου Πλοίου

ΣΝΜΜ 2019

Μάθημα 3Α: Εισαγωγή στην Μηχανική Μάθηση, Νευρωνικά Δίκτυα

Γεώργιος Παπαλάμπρου

Επίκουρος Καθηγητής ΕΜΠ

george.papalambrou@lme.ntua.gr

Εργαστήριο Ναυτικής Μηχανολογίας (Κτίριο Λ)
Σχολή Ναυπηγών Μηχανολόγων Μηχανικών
Εθνικό Μετσοβιο Πολυτεχνείο

March 18, 2019

Περιεχόμενα

Περιεχόμενο Μαθήματος

- M3A: Εισαγωγή στην Μηχανική Μάθηση, Νευρωνικά Δίκτυα. Βιβλιογραφία.
- M3B: ΝΔ: Υπολογιστικά εργαλεία
- M3Γ: Εφαρμογή ΝΔ σε μηχανές εσωτερικής καύσης/virtual sensors

Πηγές

Βασικές πηγές:

- Pattern Recognition and Machine Learning, Christopher Bishop, Springer, 2006. [Ch. 5]
[Διαθέσιμο pdf στο: https://www.microsoft.com/en-us/research/publication/pattern-recognition-machine-learning/]
- Machine Learning, Tom Mitchell, McGraw Hill, 1997. [Ch. 4]
- Neural Networks: A Comprehensive Foundation, Simon Haykin, Macmillan, 1994. [Ch. 1,2,4]
- Course MIT 6.S191: Introduction to Deep Learning, www: introtodeeplearning.com, [Slides Lesson 1]

C. Bishop, MS, UK

Pattern Recognition and Machine Learning

Christopher Bishop
Published by Springer | January 2006
View Publication

Download BibTex

This leading textbook provides a comprehensive introduction to the fields of pattern recognition and machine learning. It is aimed at advanced undergraduates or first-year PhD students, as well as researchers and practitioners. No previous knowledge of pattern recognition or machine learning concepts is assumed. This is the first machine learning textbook to include a comprehensive coverage of recent developments such as probabilistic graphical models and deterministic inference methods, and to emphasize a modern Bayesian perspective. It is suitable for courses on machine learning, statistics, computer science, signal processing, computer vision, data mining, and bioinformatics. This hard cover book has 738 pages in full colour, and there are 431 graded exercises.

Solutions for these exercises and extensive support for course instructors are provided on Christopher Bishop's page.

Now available to download in full as a PDF.

Download PDF

Research Areas

Artificial intelligence

T. Mitchell, CMU, USA

Machine Learning, Tom Mitchell, McGraw Hill, 1997.

Machine Learning is the study of computer algorithms that improve automatically through experience. Applications range from datamining programs that discover general rules in large data sets, to information filtering systems that automatically learn users' interests.

This book provides a single source introduction to the field. It is written for advanced undergraduate and graduate students, and for developers and researchers in the field. No prior background in artificial intelligence or statistics is assumed.

Chapter Outline: (or see the detailed table of contents (postscript))

G. Hinton, UTor, Ca

Geoffrey Hinton

From Wikipedia, the free encyclopedia

Geoffrey Everest Hinton CC FRS FRSC[11] (born 6 December 1947) is an English Canadian cognitive psychologist and computer scientist, most noted for his work on artificial neural networks. Since 2013 he divides his time working for Google (Google Brain) and the University of Toronto.[12][13]

With David E. Rumelhart and Ronald J. Williams, Hinton was co-author of a highly cited paper that popularized the backpropagation algorithm for training multi-layer neural networks[14]. He is viewed by some as a leading figure in the deep learning community and is referred to by some as the "Godfather of Deep Learning".[15][16][17][18][19] The dramatic image-recognition milestone of the AlexNet designed by his student Alex Krizhevsky[20] for the Imagenet challenge 2012[21] helped to revolutionize the field of computer vision.[22]

Geoffrey Hinton
FRS FRSC CC

ΝΔ - Η αρχική ιδέα: Βιολογία

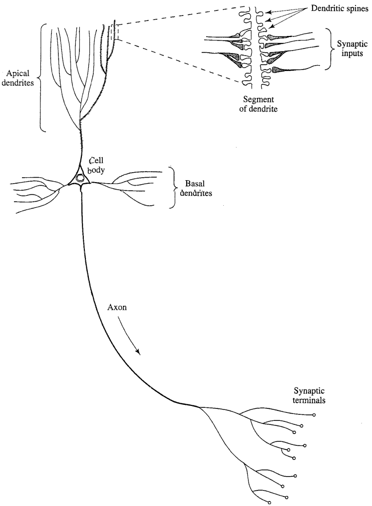

FIGURE 1.2 The pyramidal cell.

Το νευρωνικό δίκτυο

Figure 5.1 Network diagram for the two-layer neural network corresponding to (5.7). The input, hidden, and output variables are represented by nodes, and the weight parameters are represented by links between the nodes, in which the bias parameters are denoted by links coming from additional input and hidden variables x_0 and z_0. Arrows denote the direction of information flow through the network during forward propagation.

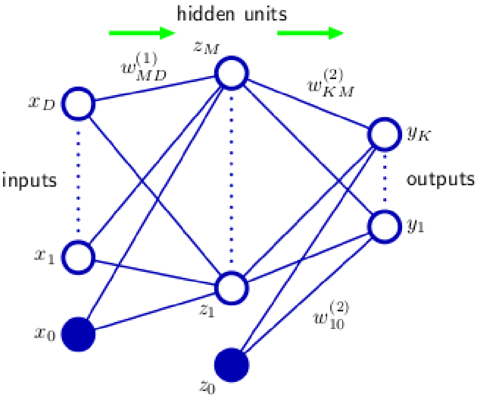

[Pattern Recognition and Machine Learning, Bishop]

Το νευρωνικό δίκτυο-Ιστορία

The term ‘neural network’ has its origins in attempts to find mathematical representations of information processing in biological systems (Widrow and Hoff, 1960; Rosenblatt, 1962; Rumelhart et al., 1986).
Indeed, it has been used very broadly to cover a wide range of different models, many of which have been the subject of exaggerated claims regarding their biological plausibility. From the perspective of practical applications of pattern recognition, however, biological realism would impose entirely unnecessary constraints.
Our focus is on neural networks as efficient models for statistical pattern recognition. In particular, we shall restrict our attention to the specific class of neural networks that have proven to be of greatest practical value, namely the multilayer perceptron.

[Pattern Recognition and Machine Learning, Bishop, Προσαρμογή από Κεφ. 5]

ALVINN ’93, [Machine Learning, Mitchell]

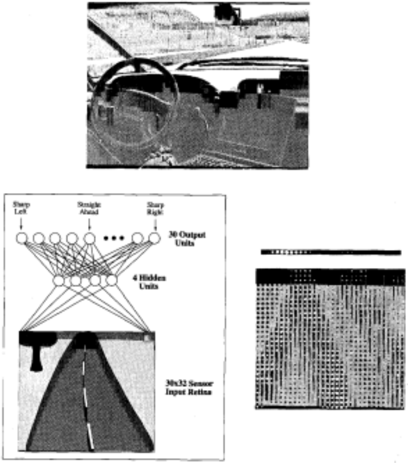

FIGURE 4.1
Neural network learning to steer an autonomous vehicle. The ALVINN system uses BACKPROPAGATION to learn to steer an autonomous vehicle (photo at top) driving at speeds up to 70 miles per hour. The diagram on the left shows how the image of a forward-mounted camera is mapped to 960 neural network inputs, which are fed forward to 4 hidden units, connected to 30 output units. Network outputs encode the commanded steering direction. The figure on the right shows weight values for one of the hidden units in this network. The 30×32 weights into the hidden unit are displayed in the large matrix, with white blocks indicating positive and black indicating negative weights. The weights from this hidden unit to the 30 output units are depicted by the smaller rectangular block directly above the large block. As can be seen from these output weights, activation of this particular hidden unit encourages a turn toward the left.

Η Μηχανική Μάθηση

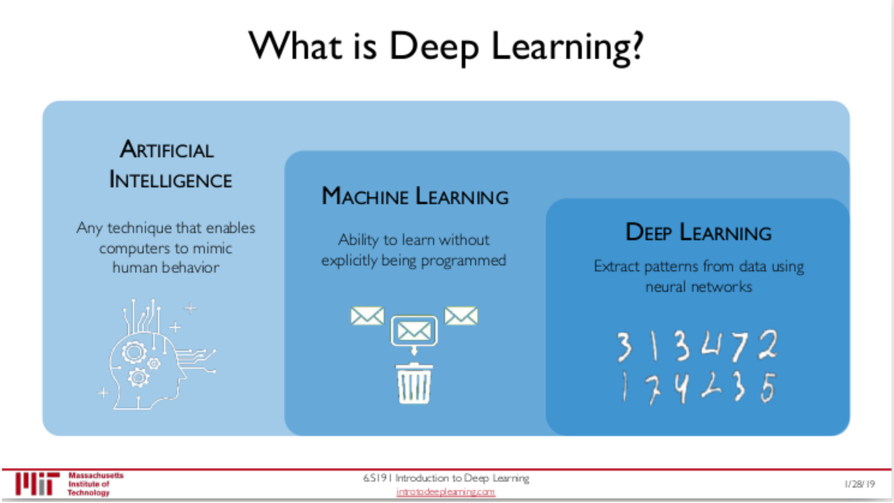

Η Επιστροφή των ΝΔ

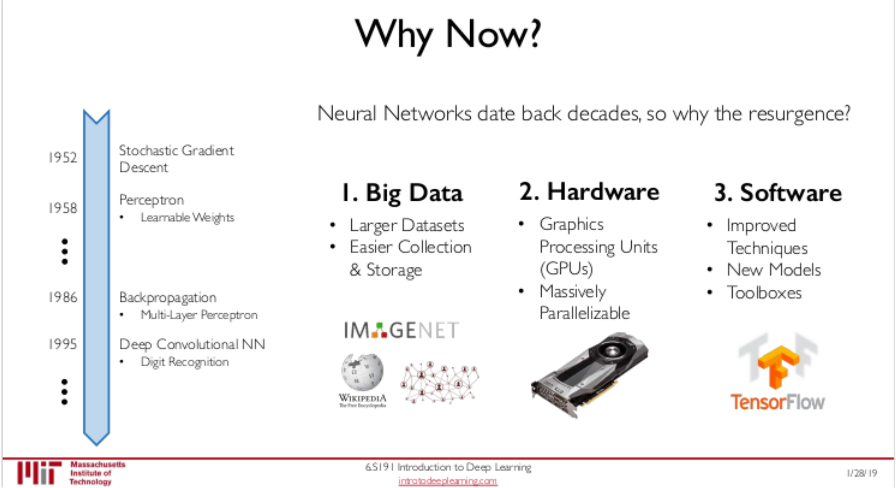

ΝΔ-Η βασική δομή

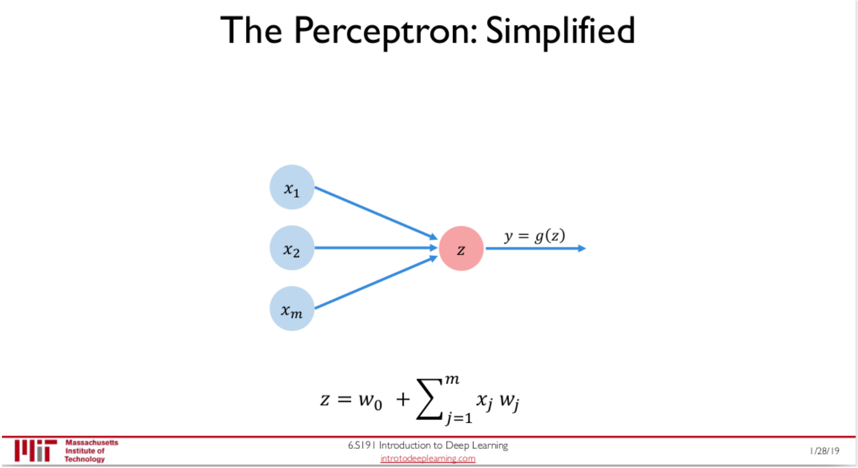

ΝΔ-Η βασική δομή

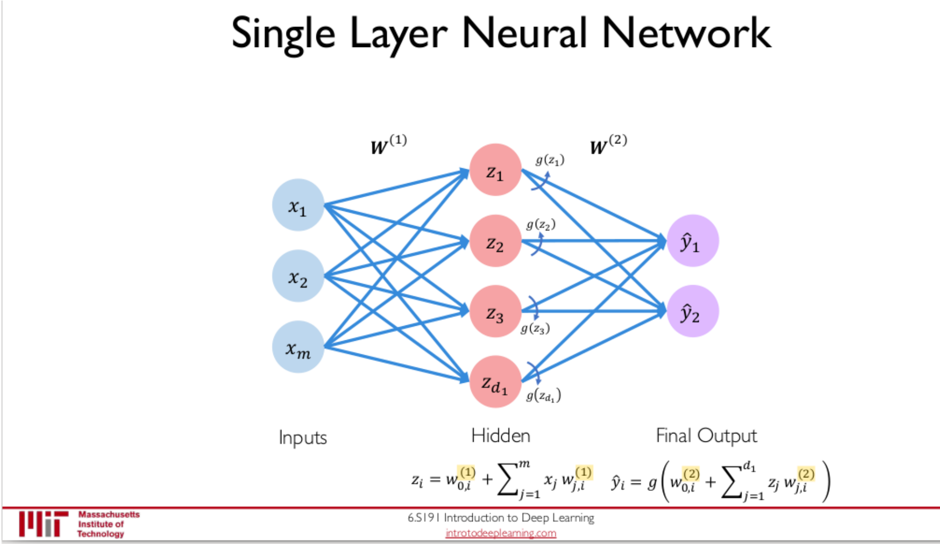

ΝΔ-Forward Propagation

The Perceptron: Forward Propagation

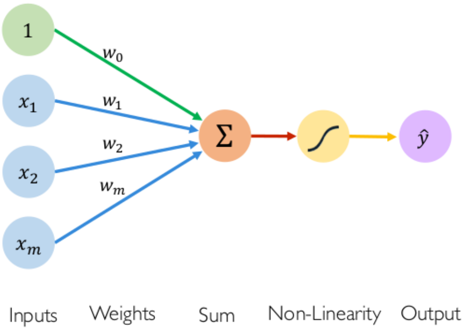

Activation Functions

$$\hat{y} = g\left(w_0 + \boldsymbol{X}^T \boldsymbol{W} \right)$$

- Example: sigmoid function

$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$

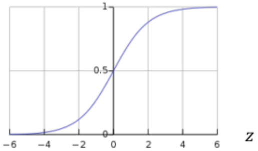

Massachusetts Institute of Technology

6.S191 Introduction to Deep Learning
introtodeeplearning.com

ΝΔ-Activation

Importance of Activation Functions

The purpose of activation functions is to ***introduce non-linearities*** *into the network*

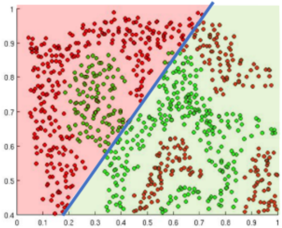

Linear Activation functions produce linear decisions no matter the network size

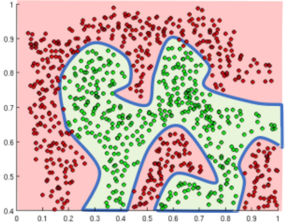

Non-linearities allow us to approximate arbitrarily complex functions

ΝΔ-Αρχιτεκτονική: Multilayer

2. Multilayer Feedforward Networks

The second class of a feedforward neural network distinguishes itself by the presence of one or more *hidden layers*, whose computation nodes are correspondingly called *hidden neurons* or *hidden units*. The function of hidden neurons is to intervene between the external input and the network output in some useful manner. By adding one or

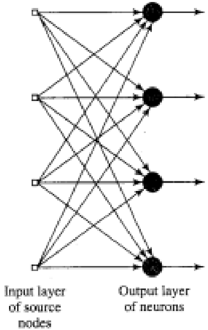

FIGURE 1.15 Feedforward or acyclic network with a single layer of neurons.

Neural Networks: A Comprehensive Foundation, Simon Haykin

ΝΔ-Αρχιτεκτονική: Hidden

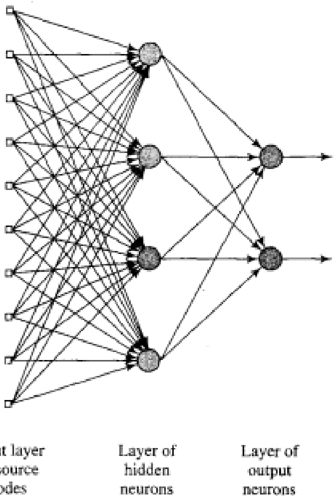

FIGURE 1.16 Fully connected feedforward or acyclic network with one hidden layer and one output layer.

Neural Networks: A Comprehensive Foundation, Simon Haykin

ΝΔ-Αρχιτεκτονική: Recurrent

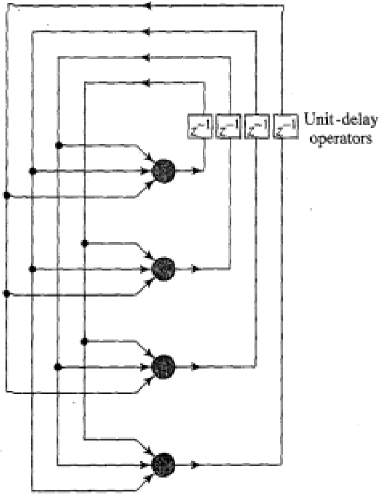

FIGURE 1.17 Recurrent network with no self-feedback loops and no hidden neurons.

Neural Networks: A Comprehensive Foundation, Simon Haykin

ΝΔ-Αρχιτεκτονική: Combined

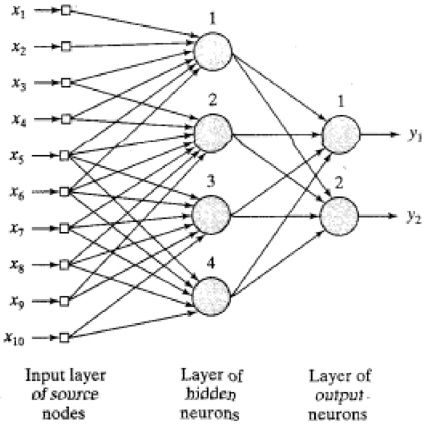

FIGURE 1.20 Illustrating the combined use of a receptive field and weight-sharing. All four hidden neurons share the same set of weights for their synaptic connections.

Neural Networks: A Comprehensive Foundation, Simon Haykin

NΔ-Learning

2.1 INTRODUCTION

The property that is of primary significance for a neural network is the ability of the network to *learn* from its environment, and to *improve* its performance through learning. The improvement in performance takes place over time in accordance with some prescribed measure. A neural network learns about its environment through an interactive process of adjustments applied to its synaptic weights and bias levels. Ideally, the network becomes more knowledgeable about its environment after each iteration of the learning process.

There are too many activities associated with the notion of "learning" to justify defining it in a precise manner. Moreover, the process of learning is a matter of viewpoint, which makes it all the more difficult to agree on a precise definition of the term. For example, learning as viewed by a psychologist is quite different from learning in a classroom sense. Recognizing that our particular interest is in neural networks, we use a definition of learning that is adapted from Mendel and McClaren (1970).

We define learning in the context of neural networks as:

> *Learning is a process by which the free parameters of a neural network are adapted through a process of stimulation by the environment in which the network is embedded. The type of learning is determined by the manner in which the parameter changes take place.*

Neural Networks: A Comprehensive Foundation, Simon Haykin

NΔ-Learning

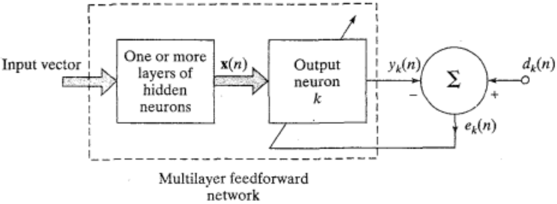

(a) Block diagram of a neural network, highlighting the only neuron in the output layer

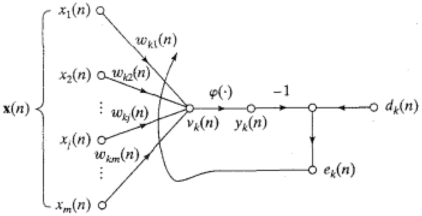

(b) Signal-flow graph of output neuron

FIGURE 2.1 Illustrating error-correction learning.

NΔ-Learning with Teacher

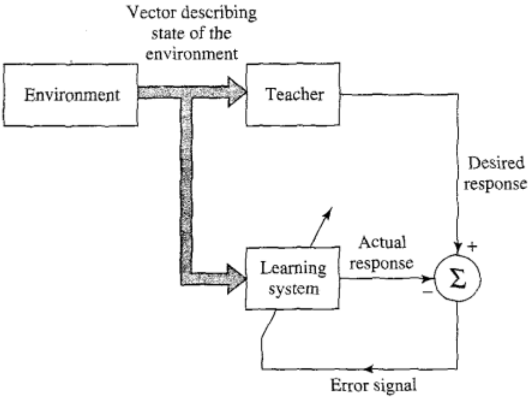

FIGURE 2.6 Block diagram of learning with a teacher.

Neural Networks: A Comprehensive Foundation, Simon Haykin

ΝΔ-Identification

- *System identification.* Let Eq. (2.19) describe the input–output relation of an unknown memoryless *multiple input-multiple output (MIMO) system*; by a "memoryless" system we mean a system that is time invariant. We may then use the set of labeled examples in Eq. (2.20) to train a neural network as a model of the system. Let $\mathbf{y}_i$ denote the output of the neural network produced in response to an input vector $\mathbf{x}_i$. The difference between $\mathbf{d}_i$ (associated with $\mathbf{x}_i$) and the network output $\mathbf{y}_i$ provides the error signal vector $\mathbf{e}_i$, as depicted in Fig. 2.11. This error signal is in turn used to adjust the free parameters of the network to minimize the squared difference between the outputs of the unknown system and the neural network in a statistical sense, and is computed over the entire training set.
- *Inverse system.* Suppose next we are given a known memoryless MIMO system whose input–output relation is described by Eq. (2.19). The requirement in this

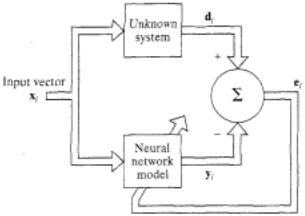

FIGURE 2.11 Block diagram of system identification.

Neural Networks: A Comprehensive Foundation, Simon Haykin

Ένα πρόβλημα: Overfitting

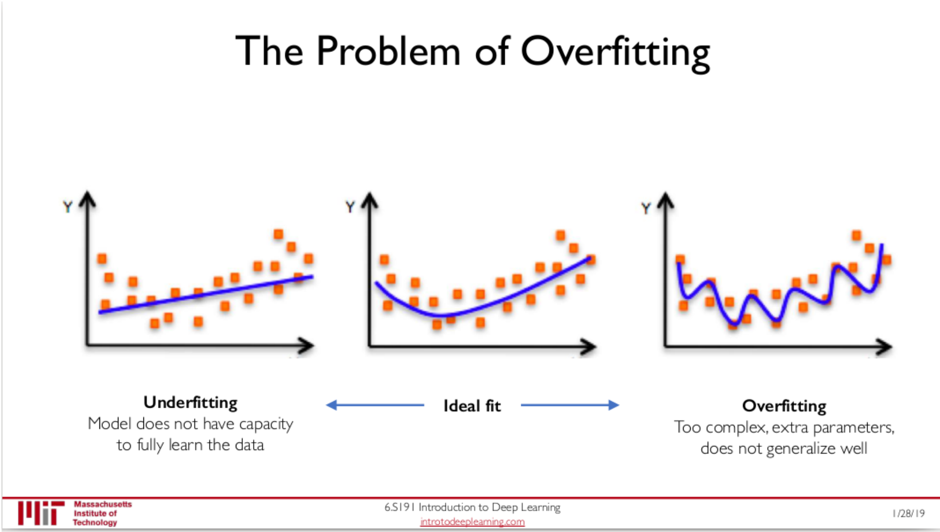

NL Classification, [Machine Learning, Mitchell]

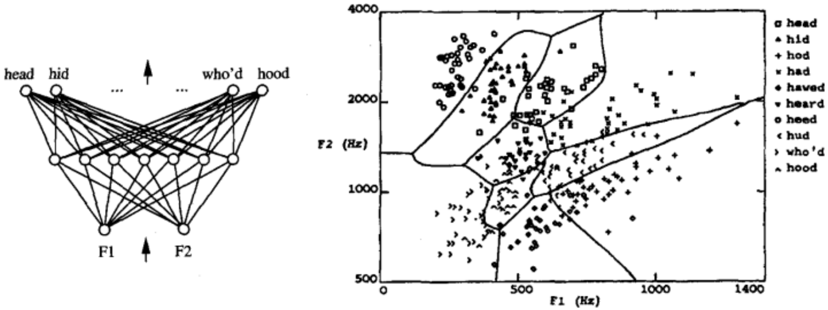

FIGURE 4.5
Decision regions of a multilayer feedforward network. The network shown here was trained to recognize 1 of 10 vowel sounds occurring in the context "h_d" (e.g., "had," "hid"). The network input consists of two parameters, F1 and F2, obtained from a spectral analysis of the sound. The 10 network outputs correspond to the 10 possible vowel sounds. The network prediction is the output whose value is highest. The plot on the right illustrates the highly nonlinear decision surface represented by the learned network. Points shown on the plot are test examples distinct from the examples used to train the network. (Reprinted by permission from Haung and Lippmann (1988).)

Ένας Virtual Sensor

Development of recurrent neural networks for virtual sensing of NOx emissions in internal combustion engines

Ivan Arsie, Cesare Pianese, Marco Sorrentino
Department of Mechanical Engineering, University of Salerno

ABSTRACT

The paper focuses on the experimental identification and validation of recurrent neural networks (RNN) for virtual sensing of NO emissions in internal combustion engines (ICE). Suited training procedures and experimental tests are proposed to improve RNN precision and generalization in predicting NO formation dynamics. The reference Spark Ignition (SI) engine was tested by means of an integrated system of hardware and software tools for engine test automation and control strategies prototyping. A fast response analyzer was used to measure NO emissions at the exhaust valve. The accuracy of the developed RNN model is assessed by comparing simulated and experimental trajectories for a wide range of operating scenarios. The results evidence that RNN-based virtual NO sensor will offer significant opportunities for implementing on-board feedforward and feedback control strategies aimed at improving the performance of after-treatment devices.

INTRODUCTION

Automotive engines and control systems are more and more sophisticated due to increasingly restrictive environmental regulations. Particularly in Diesel and GDI engines, complex after-treatment devices, such as SCR and DPF, have been introduced to meet the imposed NOx and particulate emission limits. Furthermore, on-board diagnostics for both Diesel and spark ignition (SI) engines is becoming tighter and tighter; in fact engine faults could lead to performance degradation, which may cause higher emissions. More recently, multi-fuel powertrains control units are required to identify the actual fuel composition and to adapt the strategies accordingly.

To cope with the new issues associated with novel complex hardware and to improve both powertrain performance and after-treatment efficiency, engine design methods and control and diagnostic strategies have to be revised. A key role in achieving enhanced performance is played by on-board implementation of dedicated models, either as predictors or embedded in model-based controllers or in diagnostics schemes. Virtual sensors (VSs) can be used as substitute of real sensors providing useful information about the actual process. In this scenario, virtual sensors are an interesting opportunity in achieving more challenging control and diagnostics targets.

In the last decade virtual sensors have been developed for many applications e.g. aerospace, biotechnological systems, environmental monitoring systems, to mention just few cases (Goodwin, 1999). VSs can be especially useful when: i) the direct measurements are impossible to perform, such as highly hazardous chemical or nuclear plants, (Sevilla and Pulido, 1998); ii) the available sensor might not guarantee the required sensing characteristics (i.e. accuracy, dynamic performance); iii) the sensor is too expensive or iv) does not fit within the sensing location. It is also evident that VSs can be used as feedback signal generators and their use provides opportunities to implement innovative control schemes and diagnostic. Actually, VSs must be designed in such a way as to simulate time-dependent processes and guarantee accuracy, stability and short computational time. In this work it will be shown that the proposed virtual sensor also allows a limited recourse to experiments, thus reducing costs and developing time.

Internal combustion engines are complex systems to be controlled with great precision and monitored continuously; moreover, their dynamic behavior requires both high sampling rates and rapid actuation. Therefore, on-board computational applications (e.g. models, controllers, monitoring and diagnostics algorithms) must be fast enough to ensure the required performance. VSs may support the development of new control and diagnostic concepts in achieving more challenging engine targets. Emission monitoring and control are the most complex tasks to be accomplished in an engine; this is mainly due to the complexity of simulating a process that entails advanced thermo-chemistry and fluid-dynamics within models and controllers. The state of the art mainly relies on burdensome applications where the process knowledge is stored in a large experimental database implemented on maps. A successful attempt to develop a steady-state first

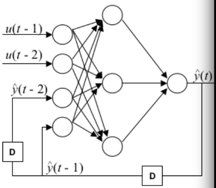

Fig. 1 – Structure of NOE RNN with variable, one output, one hidden layer output delays *D*.

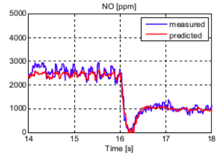

Fig. 11. Trajectories of measured and predicted NO emissions for the 2[nd] validation data set, focus on the time window (14-18 s).

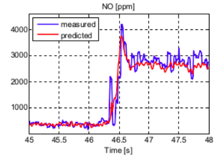

Fig. 12. Trajectories of measured and predicted NO emissions for the 2[nd] validation data set, focus on the time window (45-48 s).

ΝΔ σε θαλάσσιο περιβάλλον

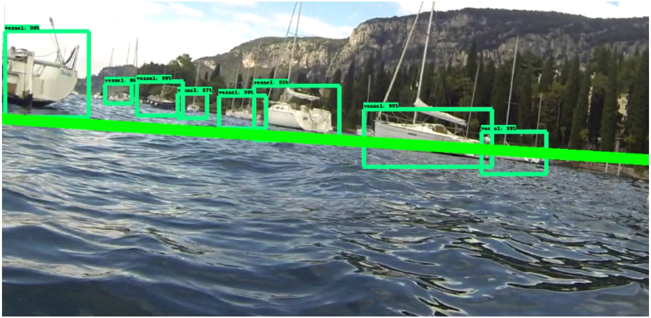

Διπλωματική Εργασία Κ. Βασιλόπουλου, ΣΝΜΜ 2018.

TO DO !

Μελετήστε από το:
Neural Networks: A Comprehensive Foundation, Simon Haykin, Macmillan, 1994
το κέφ. 1.