

Recursive Lazy Learning for Modeling and Control

Gianluca Bontempi, Mauro Birattari, and Hugues Bersini

Iridia - CP 194/6
 Université Libre de Bruxelles
 50, av. Franklin Roosevelt
 1050 Bruxelles - Belgium
 email: {gbornte,mbiro,bersini}@ulb.ac.be
<http://iridia.ulb.ac.be>

Abstract. This paper presents a local method for modeling and control of non-linear dynamical systems from input-output data. The proposed methodology couples a local model identification inspired by the lazy learning technique, with a control strategy based on linear optimal control theory. The local modeling procedure uses a query-based approach to select the best model configuration by assessing and comparing different alternatives. A new recursive technique for local model identification and validation is presented, together with an enhanced statistical method for model selection. The control method combines the linearization provided by the local learning techniques with optimal linear control theory, to control non-linear systems in configurations which are far from equilibrium. Simulations of the identification of a non-linear benchmark model and of the control of a complex non-linear system (the bioreactor) are presented. The experimental results show that the approach can obtain better performance than neural networks in identification and control, even using smaller training data sets.

1 Introduction

In this paper we present a local method to model and control an unknown dynamical system from input-output data. The idea of local approximators as alternative to global models originated in non-parametric statistics [13,11] to be later rediscovered and developed in the machine learning field [1,8]. Recent work on lazy learning (also known as memory-based or instance-based learning) gave a new impetus to the adoption of local techniques for modeling [3] and control problem [4].

Our approach extends the idea of local learning in several directions. First, we propose a model identification methodology based on the use of an iterative optimization procedure to select the best local model among a set of different candidates. Modeling a non-linear mapping using observations, requires the data analyst to make several choices involving the set of relevant variables and observations, the model structure, the learning algorithm, and the validation protocol. Our method defers all of these decisions until a prediction or a local description is requested (query-based approach). In classical methods the many options of a

local model are designed according to heuristic criteria and a priori assumptions. Here we propose an automatic procedure which searches for the optimal local model configuration, by returning for each candidate model its parameters and a statistical description of its generalization properties. As a result, a different model tuning is performed locally for each query. This idea is already presented as *local tuning* in Atkeson *et al.* [9] but no reference to any existing application is given. To the authors' knowledge, this is the first work where this approach is used to this larger extent.

The second contribution of the paper is the introduction of a new algorithm to estimate in a recursive way the model performance in cross-validation. Myers [21] introduced the PRESS statistic which is a simple, well-founded, and economical way to perform leave-one-out cross validation [12] and to assess the performance in generalization of local linear models. Here we propose a technique based on recursive least squares methods to compute the PRESS in an incremental way. Moreover, a powerful and well-founded statistical test is used to compare the performance of two alternative candidates on the basis of their cross-validation error sampling distributions.

The third contribution of the paper is a non-linear control design technique, which extensively uses analysis and design tools imported from linear control. The idea of employing linear techniques in a non-linear setting is not new in the control literature but recently had a renewed popularity thanks to methods for combining multiple estimators and controllers in different operating regimes of the system [20]. Gain scheduling [23], fuzzy inference systems [26], and local model networks [17] are well-known examples of control techniques for non-linear systems inspired by linear control. However, two strong assumptions underlie linearization control methods: an analytical description of the locus of equilibrium points is available, and the system is supposed to evolve in a sufficiently restricted neighborhood of the desired regime. Here we propose an indirect control method for performing finite-time horizon control which requires only input-output data from the observed system behavior. The approach is an example of differential dynamic programming algorithm [16,2] where the gradient computation is performed by the lazy algorithm. The controller is designed with optimal control techniques parameterized with the values returned by the linear local estimator. We show that a combination of a local estimator with a time varying optimal control can take into account the non-linearity of a system over a wider range than conventional linearized quadratic regulators (LQR).

2 Local modeling as an optimization problem

Modeling from data involves integrating human insight with learning techniques. In many real cases, the analyst faces a situation where a set of data is available, and an accurate prediction is required. Often, information about the order, the structure, or the set of relevant variables is missing or not reliable. The process of learning consists of a trial and error procedure during which the model is properly tuned on the available data. In the lazy learning approach, the estimation of

the value of the unknown function is solved giving the whole attention to the region surrounding the point where the estimation is required. The classical non-adaptive memory-based procedure essentially consists of these steps:

1. for each query point x_q , define a set of neighbors, each weighted according to some relevance criterion (e.g. the distance)
2. choose a regression function f in a restricted family of parametric functions
3. compute the regression value $f(x_q)$.

The data analyst who adopts a local regression approach, has to make a set of decisions related to the model (e.g. the number of neighbors, the weight function, the parametric family, the fitting criterion to estimate the parameters). We extend the classical approach with a method that automatically selects the adequate configuration. To do this, we simply import tools and techniques from the field of linear statistical analysis. The most important of these tools is the PRESS statistic [21], which is a simple, well-founded and economical way to perform *leave-one-out* cross-validation [12] and therefore to assess the generalization performance of local linear models. Due to its short computation time which allows its intensive use, it is the key element of our approach to modeling data. In fact, if PRESS can assign a quantitative performance to each linear model, alternative models with different configurations can be tested and compared in order to select the best one. This same selection strategy is used to select the training subset among the neighbors, as well as various structural aspects like the features to consider and the degree of the polynomial used as a local approximator [6]. The general ideas of the approach can be summarized in the following way.

1. The task of learning an input-output mapping is decomposed in a series of linear estimation problems
2. Each single estimation is treated as an optimization problem in the space of alternative model configurations
3. The estimation ability of each alternative model is assessed by the cross-validation performance computed using the PRESS statistic.

The core operation of the algorithm consists in assessing and comparing local models having different configurations. Each assessment requires the PRESS computation and a comparison with the other candidate models. In order to make these operations more effective we propose two innovative algorithms in the lazy learning method:

1. a recursive algorithm for the parametric estimation and the cross-validation of each local model. This method avoids having to restart each model evaluation from scratch and noticeably decreases the computational cost
2. a more rigorous statistical test to compare the performance of two alternative candidate models. The test does not consider only the average values of the cross-validation errors but also their sampling distributions.

2.1 The PRESS statistic and the recursive method

To illustrate the local regression procedure, we will first define some notation. Let us consider an unknown input-output mapping $f: \mathbb{R}^d \rightarrow \mathbb{R}$ of which we are

given a set of N samples $\{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$. These examples can be collected in an input matrix X of dimensionality $[N \times d]$, and in an output vector y of dimensionality $[N \times 1]$.

Given a specific query point x_q , the prediction of the value $y_q = f(x_q)$ is computed as follows. First, for each sample (x_i, y_i) , a weight w_i is computed as a function of the distance $d(x_i, x_q)$ from the query point x_q to the point x_i . Each row of X and y is then multiplied by the corresponding weight creating the variables $Z = WX$ and $v = Wy$, with W diagonal matrix having diagonal elements $W_i = w_i$. Finally, a local model is fitted solving the equation $(Z^T Z)\beta = Z^T v$ and the prediction of the value $f(x_q)$ is obtained by evaluating the model at the query point $\hat{y}_q = x_q^T (Z^T Z)^{-1} Z^T v$.

An important aspect of local learning is that along with the prediction and with the model parameters, an assessment of the performance can be easily computed. We will focus on the *leave-one-out* cross-validation procedure [12], which returns a reliable estimation of the prediction error in x_q . We define the i -th *leave-one-out* error $e^{cv}(i)$ as the difference between y_i and the prediction given by the local model centered at x_q and fitted using all the examples available except the i -th. Hence, an estimation of the prediction error in x_q is given by the average of the errors $e^{cv}(i)$ each weighted according to the respective distance $d(x_i, x_q)$. When considering a local linear model, the *leave-one-out* cross-validation can be performed without recalculating the regression parameter for each excluded example thanks to the local version of the PRESS statistic [3]:

$$\text{MSE}^{cv}(x_q) = \frac{1}{\sum_i w_i^2} \sum_i (w_i e^{cv}(i))^2. \quad (1)$$

In our modeling procedure the performance of a model in cross-validation is the criterion adopted to choose the best local model configuration. One of the most important parameters to be tuned in a local model configuration is the size of the region surrounding x_q in which the function f can be conveniently approximated by a linear local model. Such a parameter can be defined by the number of training examples which fall into the region of linearity. Consequently the task of identifying the region of linearity can be reduced to the task of finding, among the examples available, the number n of neighbors of x_q to be used in the local regression fit. Thus, we consider different models, each fitted on a different number of examples, and we use the *leave-one-out* cross-validation to compare them and to select the one for which the predicted error is smaller.

To make the procedure faster, and to avoid repeating for each model the parameter and the PRESS computation, we adopt an incremental approach based on recursive linear techniques. Recursive algorithms have been developed for model identification and adaptive control [15] to identify a linear model when data are not available from the beginning but are observed sequentially. Here we employ these methods to obtain the parameters of the model fitted on n nearest neighbors by updating the parameters of the model with $n - 1$ examples. Also, the *leave-one-out* errors $e^{cv}(i)$ are obtained exploiting partial results from the least square method and do not require additional computational overhead. Once

adopted as the weighting kernel the indicator function which assigns $w_i = 1$ to the examples used to fit the model, the recursive lazy algorithm is described by the following equations:

$$\begin{cases} P(n+1) = P(n) - \frac{P(n)x(n+1)x(n+1)^T P(n)}{1+x(n+1)^T P(n)x(n+1)} \\ \gamma(n+1) = P(n+1)x(n+1) \\ e(n+1) = y(n+1) - x(n+1)^T \beta(n) \\ \beta(n+1) = \beta(n) + \gamma(n+1)e(n+1) \\ e_{n+1}^{cv}(\hat{i}) = \frac{y(\hat{i}) - x(\hat{i})^T \beta(n+1)}{1+x(\hat{i})^T P(n+1)x(\hat{i})} \end{cases} \quad (2)$$

where, $(x(n+1), y(n+1))$ is the $n+1$ -th nearest neighbor of the query point, $P(n)$ is the recursive approximation of the matrix $(Z^T Z)^{-1}$, $\beta(n)$ denotes the optimal least squares parameters of the model fitted on the n nearest neighbors, and $e_n^{cv}(\hat{i})$, with $1 \leq i \leq n$, is the vector E_n^{cv} of *leave-one-out* errors. Once this vector is available, the formula (1) is easily computed. This value is a weighted average of the cross-validated errors and is the simplest statistic that can be used to describe the performance of the model defined by n neighbors. However, the problem of assessing the right dimension of the linearity region using a finite number of samples affected by noise requires a more powerful statistical procedure. In the following section, we will discuss in detail the method used in our model selection procedure.

2.2 The statistical test for model selection

The recursive method described in the previous section returns for each size n of the neighborhood a vector E_n^{cv} of *leave-one-out* errors. In order to select the best model, our procedure consists in increasing the number of neighbors considered when identifying the local model, until the model performance deteriorates and a departure from the region of local linearity is detected. This requires a statistical test to evaluate when the enlarged model is significantly worse than those already considered. In terms of hypothesis testing, we formulate the null hypothesis H_0 that E_n^{cv} and E_{n+1}^{cv} belong to the same distribution. To evaluate this hypothesis we use a permutation test [24] which does not require any assumptions about normality, homogeneity of variance, or about the shape of the underlying distribution. We adopt a paired version of the permutation algorithm because of the correlation between the two error vectors.

In simple words, the procedure consists in computing for each pair of errors the difference $d_i = e_n^{cv}(\hat{i}) - e_{n+1}^{cv}(\hat{i})$, and computing the value $D = \sum d_i$ which is assumed to be an instance of the random variable D^* . The sampling distribution of D^* is found by a randomization procedure [10], a computer-intensive statistical method to derive the sampling distribution of a statistic by simulating the process of sample extraction. In the permutation test, this is done by creating a high number of pseudo-samples D^b , with $b = 1, \dots, B$, derived from the actual sample D by substituting randomly a difference d_i with $-d_i$. Once the sampling

distribution of D^* is generated, a one-tailed test determines whether the null hypothesis has to be rejected.

The randomization test shows one of the main advantages of a local modeling procedure: with low computational effort it is possible to return along with the prediction and the linear local parameters also a statistical description of the uncertainty affecting these results. This property can result useful both for prediction and for control problems.

3 Lazy learning optimal control

Although non-linearity characterizes most real control problems, methods for analysis and control design are considerably more powerful and theoretically founded for linear systems than for non-linear ones. Here we propose a hybrid architecture for the indirect control of non linear discrete time plants from their observed input-output behavior. This approach combines the local learning identification procedure described in the previous section with control techniques borrowed from conventional linear optimal control

Consider a class of discrete time dynamic systems whose equations of motion can be expressed in the form

$$y(k) = f(y(k-1), \dots, y(k-ny), u(k-d), \dots, u(k-d-nu), e(k-1), \dots, e(k-ne)) + e(k), \quad (3)$$

where $y(k)$ is the system output, $u(k)$ the input, $e(k)$ is a zero-mean disturbance term, d is the relative degree and $f(\cdot)$ is some non linear function. This model is known as the NARMAX model [18]. Let us assume we have no physical description of the function f but a set of pairs $[u(k), y(k)]$ from the observed input-output behavior. Defining the information vector

$$\varphi(k-1) = [y(k-1), \dots, y(k-ny), u(k-d), \dots, u(k-d-nu), e(k-1), \dots, e(k-ne)], \quad (4)$$

the system (3) can be written in the input-output form $y(k) = f(\varphi(k-1)) + e(k)$. Consider the optimal control problem of a non linear system over a finite horizon time. Using a quadratic cost function, the solution to an optimal control problem is the control sequence U that minimizes

$$J = \frac{1}{2} y(t_f)^T P(t_f) y(t_f) + \frac{1}{2} \sum_k |y(k)^T u(k)^T| \frac{Q_k M_k}{M_k^T R_k} |y(k) u(k)| \quad (5)$$

with Q_k, M_k, R_k, P_f weighting terms designed a priori. While analytic results are not available for a generic non linear configuration, optimal control theory [25] provides the solution for the linear case. Hence, we will now present the non linear problem in a linear time varying setting.

Consider the trajectory of the dynamical system once forced by an input sequence $U = [u(1), u(2), \dots, u(t_f)]$. Assume that the system can be linearized about each state of the trajectory. Neglecting the residual errors due to the first order Taylor series approximation, the behavior of the linear system along a generic trajectory is the behavior of a linear time varying system whose state equations can be written in the form

$$\begin{aligned} y(k+1) &= A(\varphi(k))y(k) + B(\varphi(k))u(k) + K(\varphi(k)) \\ &= A_k y(k) + B_k u(k) + K_k \end{aligned} \quad (6)$$

with A_k, B_k, K_k parameters of the system linearized about the query point $\varphi(k)$. K_k is an offset term that equals zero in equilibrium points. This term requires a slight modification in the linear controller formulation. However, in order to simplify the notation, in the following we will neglect the constant term.

Optimal control theory provides the solution for the linear time varying system (6). At each time step the optimal control action is

$$u(k) = -(R_k + B_k^T P_{k+1} B_k)^{-1} (M_k^T + B_k^T P_{k+1} A_k) y(k) \quad (7)$$

where P_k is the solution to the backward Riccati equation.

$$\begin{aligned} P_k &= Q_k + A_k^T P_{k+1} A_k \\ &\quad - (M_k + A_k^T P_{k+1} B_k)(R_k + B_k^T P_{k+1} B_k)^{-1} (M_k^T + B_k^T P_{k+1} A_k) \end{aligned} \quad (8)$$

having as final condition

$$P(t_f) = P_f \quad (9)$$

The piecewise-constant optimal solution is obtained by solving the Euler-Lagrange equations, the three necessary and sufficient conditions for optimality when the final time is fixed.

$$0 = \frac{\partial H_k}{\partial u_k} = y_k^T M_k + u_k^T R_k + \lambda_{k+1}^T B_k \quad (10)$$

$$\lambda_k^T = -\frac{\partial H_k}{\partial y_k} = y_k^T Q_k + u_k^T M_k^T + \lambda_{k+1}^T A_k \quad (11)$$

$$\lambda_k^T = y_k^T P_f \quad (12)$$

with $\lambda_k = P_k y_k$ adjoint term in the augmented cost function (Hamiltonian)

$$H_k = J + \lambda_{k+1}^T (A_k y(k) + B_k u(k)) \quad (13)$$

The Euler-Lagrange equations do not hold for non linear systems. Anyway, if the system can be represented in the form (6), formula (10) can be used to compute the derivative of the cost function (5) with respect to a control sequence U . This is the idea of differential dynamic programming algorithms [16] which require at each time k the matrices A_k, B_k , linearizations of the system dynamics along the trajectory forced by the input sequence.

As discussed in Section 2, our modeling procedure performs system linearization with minimum effort, no a priori knowledge and only a reduced amount of data. Hence, we propose an algorithm for non linear optimal control, formulated as a gradient based optimization problem and based on the local system linearization.

The algorithm searches for the sequence of input actions

$$U^{opt} = \arg \min_{U^i} J(U^i) \quad (14)$$

that minimizes the finite-horizon cost function (5) along the future t_f steps. The cost function $J(U^i)$ for a generic sequence U^i is computed simulating forward for t_f steps the model identified by the local learning method. The gradient of $J(U^i)$ with respect to U^i is returned by (10).

These are the basic operations of the optimization procedure executed each time a control action is required:

1. forward simulation of the lazy model forced by a finite control sequence U^i of dimension t_f
2. linearization of the simulated system about the resulting trajectory
3. computation of the resulting finite cost function $J(U^i)$
4. computation of the gradient of the cost function with respect to simulated sequence
5. updating of the sequence with a gradient based algorithm.

Once the search algorithm returns the optimal solution U^{opt} , the first action of the sequence is applied to the real system (receding horizon control strategy [9]). Let us remark how the lazy learning model has a twofold role in the algorithm: (i) at step 1 it behaves as an approximator which predicts the behavior of the system once forced with a generic input sequence (ii) at step 2 it returns a linear approximation to the system dynamics.

Atkeson *et al.* [4] and Tanaka [27] applied infinite-time LQR regulator to non linear systems linearized with lazy learning and neuro-fuzzy models. The drawback of these approaches is that an equilibrium point or a reference trajectory is required. Also, they make the strong assumption that the state of the system will remain indefinitely in a neighborhood of the linearization point. As discussed above, the advantage of the proposed approach is that these requirements do not need to be satisfied. First, lazy learning is able to linearize a system in points far from equilibrium. Secondly, the time varying approach makes possible the use of a linear control strategy even though the system operates within different linear regimes.

Remark: we make the assumption that the parameters returned by the local models are a real description of the local behavior (certainty equivalence principle). This is a restricting assumption which requires a sufficient degree of accuracy in the approximation. However, we see in the optimal control theory a possible solution to this limitation. In fact, stochastic optimal control theory provides a formal solution to the problem of parameter uncertainty in control systems (dual control [14]). Further, our modeling procedure can return at no

additional cost a statistical description of the estimated parameters (see Section 2.2). Hence, future work will focus on the extension of the technique to the stochastic control case.

4 Experiments

4.1 The identification of a non-linear discrete time system

The approach has been applied to the identification of a complex non-linear benchmark proposed by Narendra and Li [22]. The discrete time equations of the system are:

$$\begin{cases} x_1(k+1) = \left(\frac{x_1(k)}{1+x_1^2(k)} + 1 \right) \sin(x_2(k)) \\ x_2(k+1) = x_2(k) \cos(x_2(k)) + x_1(k) e^{-\frac{x_2^2(k)+x_1^2(k)}{8}} \\ \quad + \frac{u^2(k)}{1+u^2(k) + 0.5 \cos(x_1(k) + x_2(k))} \\ y(k) = \frac{x_1(k)}{1+0.5 \sin(x_2(k))} + \frac{x_2(k)}{1+0.5 \sin(x_1(k))} \end{cases} \quad (15)$$

where (x_1, x_2) is the non observable state and only the input u and the output y are accessible. We model the non observable system in the input-output form $y(k+1) = f(y(k), y(k-1), y(k-2), y(k-3), u(k))$. We use an initial empty database which is updated all along the identification. We perform the identification for 1500 time steps with a test input $u(k) = \sin(\frac{2\pi k}{10}) + \sin(\frac{2\pi k}{25})$. The plot in Fig. 1a shows the model and the system output in the last 200 points, while the plot in Fig. 1b shows the identification error.




Fig. 1. Non linear system identification results: a) system (solid) and model (dotted) outputs, b) identification error

We obtain a good performance in modeling this complex non observable system. The experiment reproduces exactly the one proposed in Narendra and

Li [22], with the sole difference that in our case we used a data set of 1500 points instead of the 500,000 used in the cited paper. Unfortunately, Narendra and Li do not report any quantitative index of the neural network (4-layer feed-forward) performance and present only a plot of the results. Anyway, a qualitative comparison between Fig. 11.5 in their article and Fig. 1 of this paper, shows clearly how our method outperforms Narendra's neural identification, even using much less data.

4.2 The control of the bioreactor

Consider, as second example, the bioreactor system [19], a challenging benchmark in control for its non linearity and because small changes in parameters value can cause instability. The bioreactor is a tank containing water, nutrients, and biological cells. Nutrients and cells are introduced into the tank where the cells mix with the nutrients. The state of this process is characterized by the number of cells (c_1) and the amount of nutrients (c_2). Bioreactor equations of motion are the following:

$$\begin{cases} \frac{dc_1}{dt} = -c_1 u + c_1(1-c_2) e^{\frac{c_2}{\gamma}} \\ \frac{dc_2}{dt} = -c_2 u + c_1(1-c_2) e^{\frac{c_2}{\gamma}} \left(\frac{1+\beta}{1+\beta-c_2} \right) \end{cases} \quad (16)$$

with $\beta = 0.02$ and $\gamma = 0.48$. In our experiment the goal was to stabilize the multivariable system about the unstable state $(c_1^*, c_2^*) = (0.2107, 0.726)$ by performing a control action each 0.5 seconds.

We use the control algorithm described in Section 3. The system is modeled in the input-output form (3) having the orders $ny = 2$, $nu = 1$, $me = 0$, $d = 1$. The horizon of the control algorithm is fixed to $t_f = 5$. The initial state conditions are set by the random initialization procedure defined in [19]. We initialize the lazy learning database with a set of 1000 points collected by preliminarily exciting the system with a random uniform input. The database is then updated on-line each time a new input-output pairs is returned by the simulated system. The plot in Fig. 2a shows the output of the two controlled state variables, while the plot in Fig. 2b shows the control action.

We have better results than Bersini and Gorrini [7] who used an optimal neuro-controller (MLP network with ten neurons in the hidden layer) trained with a much higher number of samples (10,000), and applied control actions at a higher frequency (each 0.01 seconds).

Both identification and control examples show that, by using local techniques, it is possible to deal with complex systems on a wide non-linear range, with no a priori knowledge about the underlying dynamics and using less data than global estimators.

5 Conclusions and future developments

Local modeling techniques are a powerful technique for learning from limited amount of data by providing at the same time an useful insight on the local




Fig. 2. Control results: a) controlled variables (solid) and references (dotted), b) control action.

behavior of the system being modeled. Furthermore, together with the required prediction and/or parametric description, they return a statistical distribution of the uncertainty affecting this information. We proposed an innovative algorithm to improve the performance of local modeling techniques which is based on a recursive version of the cross-validation and a statistical model selection. In control literature, local controllers have generally a restricted range of operating conditions. Here, we proposed a controller, which although making extensive use of local techniques, works on an extended range of operating conditions. These characteristics makes of it a promising tool for intelligent control systems, inspired to traditional engineering methods but able to deal with complex non linear systems. Successful applications of the method to simulated identification and control problems were presented. Future developments will concern the combination of the local modeling technique with other certainty equivalence controllers (e.g. minimum variance controller, pole placement) and the extension of the method to stochastic dual control.

Acknowledgments Gianluca Bontempi was supported by the European Union TMR Grant FMBICT960692. Mauro Birattari was supported by the F.I.R.-S.T. program of the Région Wallonne, Belgium.

References

1. D.W. Aha. Incremental, instance-based learning of independent and graded concept descriptions. In *Sixth International Machine Learning Workshop*, pages 387-391, San Mateo, CA, 1989. Morgan Kaufmann.
2. C. G. Atkeson. Using local optimizers to speed up global optimization in dynamic programming. In J. D. Cowan, G. Tesauro, and J. Alsppector, editors, *Advances in Neural Information Processing Systems 6*, pages 663-670. Morgan Kaufmann, 1994.
3. C.G. Atkeson, A.W. Moore, and S. Schaal. Locally weighted learning. *Artificial Intelligence Review*, 11(1-5):11-73, 1997.
4. C.G. Atkeson, A.W. Moore, and S. Schaal. Locally weighted learning for control. *Artificial Intelligence Review*, 11(1-5):75-113, 1997.

5. J.K. Benedetti. On the non parametric estimation of regression functions. *Journal of the Royal Statistical Society, Series B*, (39):248-253, 1977.
6. H. Bersini, M. Birattari, and G. Bontempi. Adaptive memory-based regression methods. In *Proceedings of the 1998 IEEE International Joint Conference on Neural Networks*, 1998, to appear.
7. H. Bersini and V. Gorrini. A simplification of the back-propagation-through-time algorithm for optimal neurocontrol. *IEEE Trans. on Neural Networks*, 8(2):437-441, 1997.
8. L. Bottou and V.N. Vapnik. Local learning algorithms. *Neural Computation*, 4(6):888-900, 1992.
9. D.W. Clarke. *Advances in Model-Based Predictive Control*. Oxford University Press, 1994.
10. P.R. Cohen. *Empirical Methods for Artificial Intelligence*. The MIT Press, Cambridge, MA, 1995.
11. B. V. Dasarthy. *Nearest neighbor (NN) norms : NN pattern classification techniques*. IEEE Computer Society Press, 1991.
12. B. Efron and R.J. Tibshirani. *An Introduction to the Bootstrap*. Chapman and Hall, New York, NY, 1993.
13. V.A. Epanechnikov. Non parametric estimation of a multivariate probability density. *Theory of Probability and Its Applications*, (14):153-158, 1969.
14. A.A. Fe'dbaum. *Optimal control systems*. Academic Press, New York, NY, 1965.
15. G.C. Goodwin and K. S. Sin. *Adaptive Filtering Prediction and Control*. Prentice-Hall, 1984.
16. D. Jacobson and D. Mayne. *Differential Dynamic Programming*. Elsevier Sci. Publ., New York, 1970.
17. T.A. Johansen and B.A. Foss. Constructing narmax models using armax models. *International Journal of Control*, 58:1125-1153, 1993.
18. I.J. Leontaris and S.A. Billings. Input-output parametric models for non-linear systems. *International Journal of Control*, 41(2):303-344, 1985.
19. W.T. MillerIII, R.S. Sutton, and P.J. Werbos, editors. *Neural Networks for Control*. The MIT Press, 1990.
20. R. Murray-Smith and T.A. Johansen, editors. *Multiple Model Approaches to Modeling and Control*. Taylor and Francis, 1997.
21. R.H. Myers. *Classical and Modern Regression with Applications*. PWS-KENT, Boston, MA, 1990.
22. K.S. Narendra and S.M. Li. Neural networks in control systems. In M.C. Mozer Paul Smolensky and D.E. Rumelhart, editors, *Mathematical Perspectives on Neural Networks*, chapter 11, pages 347-394. Lawrence Erlbaum Associates, 1996.
23. J.S. Shamma and M. Athans. Gain scheduling: Potential hazards and possible remedies. *IEEE Control Systems Magazine*, pages 101-107, June 1992.
24. S. Siegel and Jr. N.J. Castellan. *Non Parametric Statistics for the Behavioral Sciences*. McGraw-Hill International, 2nd edition, 1988.
25. R.F. Stengel. *Stochastic optimal control: theory and application*. John Wiley and Sons, New York, NY, 1986.
26. T. Takegi and M. Sugeno. Fuzzy identification of systems and its applications to modeling and control. *IEEE Transactions on System, Man and Cybernetics*, 15(1):116-132, 1985.
27. K. Tanaka. Stability and stabilizability of fuzzy-neural-linear control systems. *IEEE Transactions on Fuzzy Systems*, 3(4):438-447, November 1995.