Z. Gao, M. Peek and P. J. Antsaklis, "Learning for the Adaptive Control of a Large Flexible Structure," Proc. of the 3rd IEEE International Symposium on Intelligent Control, pp. 508-512, Arlington, VA, Aug. 24-26, 1988.

Learning for the Adaptive Control of Large Flexible Structures

by

Zhiqiang Gao+, Mark Peek++ and Panos J. Antsaklis+

+Department of Electrical and Computer Engineering
University of Notre Dame
Notre Dame, IN 46556

++Tellabs Research Center
3702 N. Main Street
Mishawaka, IN 46545

Abstract

An important problem in the adaptive control of large flexible structures is to appropriately select the adaptive controller parameters so that good performance is obtained. There is a lack of systematic methods in selecting these parameters mainly because of the mathematical complexity in studying the transient response of nonlinear systems. In this paper a method, based on machine learning, to solve this problem is introduced and discussed. These results are then used to propose an intelligent adaptive control system where the parameters in the adaptive controller are to be tuned on-line without human supervision.

1.0 Introduction

Large flexible space structures pose unique control problems because of the complexity of their dynamic behavior, the limited knowledge of the model, the time-varying elements, and the disturbances which occur. Adaptive control appears to be an effective way to control such flexible space structures[8]. In the literature on adaptive control, the main issues are stability and convergence. Unfortunately, once the stability and convergence are guaranteed for the system, there is almost no methodology to assign the design parameters (within the range imposed by stability constraints) to optimize or simply improve the closed-loop system performance. It is desirable of course to have some analytical relations between these parameters and the performance so that, first, good performance is obtained by choosing the parameters correctly, and second, once the performance starts degrading during control execution because of some external disturbances or changes in the dynamics of the plant, it can be recovered as much as possible by adjusting the appropriate parameters. This problem seems intractable mathematically, especially for the complex large flexible space structures although some progress has been reported recently for simplified models[3,4]. In addition, the set of design parameters may itself need to adapt to environmental changes. This happens when the nominal adaptive system reaches the limits of its adaptation and cannot tolerate further environmental changes. A way to effectively address these problems is to determine and store desirable sets of parameters for different operating regions of the adaptive systems. Before we introduce this approach to the problems, note that the problem of tuning design parameters in adaptive control systems can be divided into two parts:

Part 1: Given a closed-loop adaptive control system, assign all the assignable parameters in the adaptive controller so that a predefined measurement of the system performance is optimized.

Part 2: In the same system above, find the relation of each parameter and the performance measures.

Unfortunately, very little can be found in the adaptive control literature that addresses the above problems. In practice, the control designers usually settle on tedious trial and error type of approaches. This approach is not just time consuming, but more importantly, may result in a reduced stability margin, poor performance, wasted control energy, etc., for missing the optimal set of parameters.

In this paper, a new approach based on machine learning is proposed. It is shown that Learning by Observation and Discovery can be effectively used in the adaptive control design, and in particular in optimizing the system performance. The method resembles one of the early approaches which utilizes the Hill Climbing method[11]. It is shown that an expert system guided optimization scheme is well suited for the knowledge acquisition here. The search for the optimal performance is formulated as an unconstrained nonlinear optimization problem where the variables are the parameters in the adaptive controller and the cost function is the performance index (PI) which is defined as a weighted sum of the root-square-error (MSR), the maximum error (ME) and the settling time (ST). In this proof of concept type of system, we would like to make the problem simple at the beginning by assuming that the performance surface is strictly quasiconvex; we do not consider the local minimum problem in this paper. We also assume that there is no constraint on the values of the parameters.

The learning system is built on top of the adaptive controller and it employs a knowledge-based system which consists of a rulebase and a database. The rules in the rulebase are constructed by using the Hooke and Jeeves method[2] for multidimensional search and a line search method that is explained in Section 4.2. In running the learning system, the design parameters are continually perturbed by the knowledge-based system until an optimal PI is obtained according to predetermined criteria. During this process, a sensitivity analysis of the MSR, the ME and the ST with respect to each of the parameters is carried out at each step, the results are averaged over all of the steps and placed in a table in the database. The corresponding parameters for the optimal PI are placed in the database also. This approach is mainly an off-line learning process. In addition, the results over different operating regions and various types of the disturbances can be used in building an intelligent adaptive control system where the design parameters are tuned automatically on-line to overcome severe environmental changes that can not be dealt with an ordinary adaptive control system. This is further discussed in Section 6.

The problem of learning in automatic control systems has been studied considerably in the past, especially in the late 60's, and it has been the topic of numerous papers and books[7,9,11-14]. References[7,11,13] provide surveys on the early learning techniques. All of these problems involve a process of classification, in which all or part of the a priori information required is unknown or incompletely known. The elements or patterns that are presented to the control system are collected into groups that correspond to different pattern classes or regions[13]. Thus learning was viewed as the estimation or successive approximation of the unknown quantities of a function[7]. The approaches developed for such learning problems can be separated into two categories: deterministic and stochastic. Among these methods, Hill Climbing is pertinent to the problem addressed in this paper and it is discussed in the next section. Note that the description given here are brief by necessity. More details will be presented in future publications.

The development of Artificial Intelligence (AI) theory and applications provides a solid background for further study of learning control systems. Machine learning is an important part of AI and is currently being applied to many fields[5,6,10]. Various AI software such as expert system shells, allow convenient use of AI methodologies in solving practical problems. Therefore we believe that we are in a much better position today to study the learning problem in control systems.

The learning theory is discussed briefly in Section 2 to clarify what we mean by learning and what we expect from it. In

0-8186-2012-9/89/0000/0508$01.00 © 1989 IEEE

Z. Gao, M. Peek and P. J. Antsaklis, "Learning for the Adaptive Control of a Large Flexible Structure," P roc. o f t he 3 rd I EEE I nternational S ymposium o n I ntelligent C ontrol , pp. 508-512, Arlington, VA, Aug. 24-26, 1988.

Section 3 we introduce the adaptive control system. The learning scheme, and the implementation of the learning system are discussed in Section 4. The simulation results are shown in Section 5, and the use of these results in building an intelligent adaptive control system is described in Section 6. Finally in Section 7, some concluding remarks are made.

2.0 Hill Climbing and Learning by Discovery

What does learning mean in AI? H.A. Simon[10] states that "learning denotes changes in the system that are adaptive in the sense that they enable the system to do the same task or tasks drawn from the same population more effectively the next time." However, the term "learning", like a lot of other everyday terms, is used broadly and vaguely and when it is carried over to technical fields it often causes confusion. Rather than arguing what is "learning" and what is not, our objective is to apply the learning theory in AI to improve the performance of the adaptive control systems.

In AI, learning has been classified on the basis of the underlying learning strategies[10]. One of them is the so called learning from observation and discovery (also called unsupervised learning). It is a very general form of inductive learning without an external teacher. There are two sub-classes concerning the degree of interaction with the external process. One is passive observation, where the learner classifies and taxonomizes observations of multiple aspects of the environment. Another is active experimentation, where the learner perturbs the environment to observe the results of its perturbations. As a result, a learning system may acquire rules of behavior, descriptions of physical objects, problem solving heuristics, and many other types of knowledge useful in the performance of a wide variety of tasks.

One of the early approaches utilized in learning control systems, the Hill Climbing method[11], is quite similar to the type of learning discussed above, although it has not received much attention in control. There are three successive steps in the method: i) perturb system parameters in a prescribed set of directions; ii) find the new PI, which is a result of the performance evaluation; iii) move the parameters along the gradient whenever the PI does not fall within an acceptable limit. This method is usually used in systems with stochastic perturbation. In one example[11], the performance index is the expected value of the square error in a system and the goal is to minimize it. (This is actually a search for the bottom of a "valley", rather than the top of a "hill". Still, it is usually referred as hill climbing.) Here we are perturbing the parameters systematically and we monitor the system performance, and in that our method resembles the Hill Climbing approach. Using optimization techniques and an AI learning method, namely learning by observation and discovery, we built a learning system for the adaptive control system of large flexible space structures. This is shown in Section 4.

3.0 The Adaptive Control Problem

In this paper, a learning system is introduced for the selections of parameters in the adaptive controller of a large flexible space structure. The adaptive control system to be considered is shown in Figure 1, where the plant is a large space antenna[8]. The plant is decoupled into two single-input and single-output subsystems, where each one has six modes, known as boom-dish modes. For simplicity we only consider one of them, that is a single input and single output system of order twelve. The control action is provided by

$$u_p = k\tilde{r} \tag{1}$$

where $\tilde{r}^T = [\tilde{e}_y \;\; x_m \;\; u_m]$, $u_p \in R$, $u_m \in R$, $\tilde{e}_y \in R$, $x_m \in R^r$, $k \in R^{(r+2)}$, r is the dimension of the model system. The gain matrix k has the form

$$k = k_P + k_I \tag{2}$$

where k_I and k_P are the integral and proportional gain respectively. They are given by the following equations:

$$\dot{k}_I = -\sigma_1 k_I + L e_y \tilde{r}^T T \tag{3}$$

$$\dot{k}_P = -\sigma_2 k_P + \bar{L} e_y \tilde{r}^T \bar{T} \tag{4}$$

The parameters in the set, $\{\sigma_1, \sigma_2, L, \bar{L}, T, \bar{T}\}$, are to be chosen in the design process. The choice of these parameters determines the system performance. Unfortunately, there is no systematic guides available for designers to make the best possible choice. The stability and convergence analysis could provide the stability bound on these values, but due to the complexity of the transient response analysis in this nonlinear system, the analytical relation between

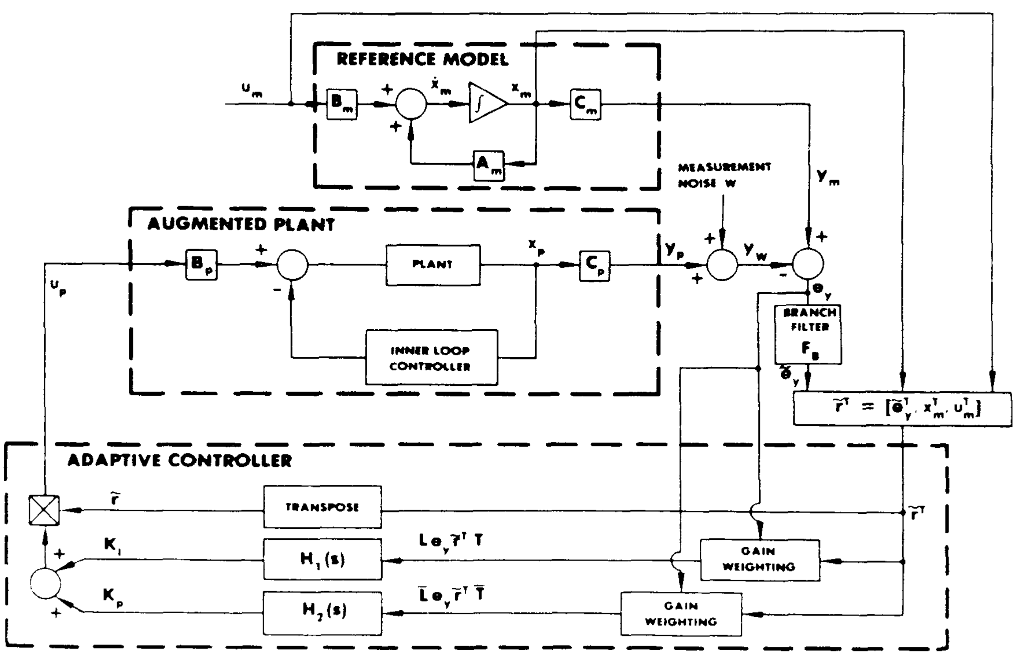

Figure 1 Adaptive control system block diagram

Z. Gao, M. Peek and P. J. Antsaklis, "Learning for the Adaptive Control of a Large Flexible Structure," Proc. of the 3rd IEEE International Symposium on Intelligent Control, pp. 508-512, Arlington, VA, Aug. 24-26, 1988.

performance and parameters is difficult to find and this area is still the subject of current research[3,4]. To reduce the problem complexity, we fix T and $\bar{T}$, which leaves only σ_1, σ_2, L and $\bar{L}$ to choose. Let S be the set $\{\sigma_1, \sigma_2, L, \bar{L}\}$ and S_{opt} denote the S that optimizes the system performance.

4.0 Learning in Adaptive Control.

The purpose of introducing learning in adaptive control system is to choose the optimal set, $S = S_{opt}$, automatically. Also we would like to obtain information on the relationship between parameters and performance, e.g., which of the components in the PI is most sensitive to the parameter L, and vice versa.

4.1 The Learning Scheme

As discussed in the introduction, the parameter choosing problem in this adaptive control system has two parts:

Part I: For certain operating region (reference model and inputs), find an optimal set $S = S_{opt}$ such that it optimizes the system performance.

Part II: Extract knowledge of the relation between parameters and performance.

To solve the problem, a method based on the learning theory in Section 2 is introduced. The learning system configuration is shown in Figure 2. The learning process involves successive off-line simulations of the adaptive control system where the knowledge-based system changes the design parameters based on the performance analysis of the previous simulations results.

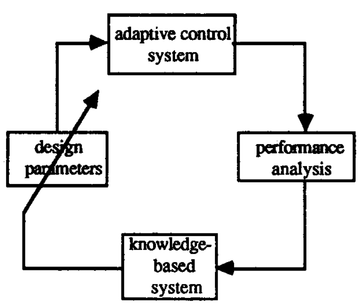

Figure 2 Learning in adaptive control system

Here the environment of the learning is the closed-loop adaptive control system. The learning by discovery strategy is applied using the active experimentation and perturbing the design parameter set, S, of the adaptive controller. The learner, which perturbs S, is a knowledge-based system with a rulebase and a database. The learning process is described as follows: Given an initial value of parameter set S, start the simulation of the adaptive control system. The knowledge-based system monitors the system performance and perturbs the parameters in S in a systematic way according to the rules in its rulebase. The rulebase is constructed by using optimization algorithms which guide the search for the valley in the performance surface since the objective is to minimize the PI. The performance analysis provides the PI as well as its components the RMS, the ME and the ST, which are the performance measures of the system. When the valley, or the minimum of the PI, is reached, the corresponding S is placed in the database as the S_{opt} and the learning process for the given operating region terminates. The system is designed so that during the process, additional knowledge is extracted, such as the sensitivity of each parameter to the performance, and stored in the database.

4.2 Implementation of the Learning System

For the implementation of the learning scheme discussed above, the search for the optimal set of parameters is formulated as an expert system guided optimization problem. The cost function to be minimized is the performance index which is defined as

$$PI = w_1 RMS + w_2 ME + w_3 ST \qquad (5)$$

where w_1, w_2 and w_3 are weighting factors. In the experiment, w_1, w_2 and w_3 are chosen as 100, 10 and 0.1 respectively. In general, the PI should include all of the important quantities for the evaluation of the system performance with the weights reflecting their relative importance together with their relative sizes. The problem of minimizing the PI is treated as an unconstrained nonlinear optimization problem. It is assumed that the performance surface is strictly quasiconvex[2] and that there are no limits on the values of the parameters. The problems beyond these assumptions will be considered in the future research.

The expert system monitors the changes in the parameters following the Hooke and Jeeves method using line search. Although the details can be found in [2], the method is outlined here for completeness. There are two types of search involved in this method, exploratory search and pattern search. The exploratory search is a line search along the coordinate axes in the parameter space. The pattern search is a line search along the pattern directions defined as $\mathbf{x}_{i+1} - \mathbf{x}_i$, where $\mathbf{x}_i$ is the point reached in exploratory search at the ith iteration. Each exploratory is followed by a pattern search. For example, given $\mathbf{x}_0$, the exploratory search along the coordinate directions produces the point $\mathbf{x}_1$. Now a pattern search along a direction $\mathbf{x}_1 - \mathbf{x}_0$ leads to the point $\mathbf{y}$. Another exploratory search starting from $\mathbf{y}$ gives the point $\mathbf{x}_2$. The next pattern search is along the direction $\mathbf{x}_2 - \mathbf{x}_1$, yielding $\mathbf{y}'$. The process is then repeated. The speed of this method was significantly better than that of the cyclic method which was used initially.

The line search is controlled by the rules which are constructed by the following strategy. Under the assumptions above, there is only one minimum along each of the coordinates in the parameter space, and the interval of uncertainty for each parameter is $(-\infty, \infty)$. Assume that the initial value of S, which are chosen by some heuristics or intuition, has no zero element, and the simulation of the closed-loop adaptive control system gives finite output. At the beginning of the line search, a parameter is either increased or decreased by 10% of its initial value. If the PI decreases, then repeat the change and run the simulation again. If the PI decreases in three consecutive steps, the direction is correct, increase the step size by a factor of 10. Continue this until the PI increases, the direction is wrong, and the minima has been passed. Go back to the last value and change the parameter with step size ten times smaller in the reversed direction. When the PI increases in both directions, decrease the step size by a factor of ten again. In the case where both changes to the PI and to the parameter fall below predetermined small numbers, the line search for this parameter stops, and the line search for the next parameter begins. This method offered good results under a variety of conditions.

Both the Hooke and Jeeves method and the above line search strategy are encoded into the rules of the rulebase in the knowledge-based system. The expert system controls the simulations of the adaptive control system and the search for an optimal set of parameters. This is essentially the Part I of the learning problem. For Part II, the knowledge acquisition in the learning process is accomplished by the sensitivity analysis at each step of the simulation controlled by the knowledge-based system. The results are then averaged over all of the steps and then placed into a look-up table, called the sensitivity table, where the columns represent the RMS, the ME and the ST, and the rows represent σ_1, σ_2, L and $\bar{L}$ respectively. The values in the table are nonnegative and show the magnitude of sensitivity of the RMS, the ME and the ST with respect to each of the parameters. This table is used for the intelligent adaptive control system in Section 6.

For practical implementation, the adaptive control algorithm in (3) and (4) is discretized and coded using the backward differencing scheme. The simulation program is written in the C language. The knowledge-based system is implemented using CxPERT, an expert system shell which is based on the C language. The advantage of this configuration is that there is no interface problem between the expert system and the numerical simulation program. The simulation runs on an IBM PC AT machine.

Z. Gao, M. Peek and P. J. Antsaklis, "Learning for the Adaptive Control of a Large Flexible Structure," P roc. o f t he 3 rd I EEE I nternational S ymposium o n I ntelligent C ontrol , pp. 508-512, Arlington, VA, Aug. 24-26, 1988.

5.0 Simulation Results

Simulations are performed on the model of the large space antenna. Only the boom-dish modes are considered and only one subsystem is used. The inner loop design may be different from the one in [8] and there is no limit on the magnitude of the input to the plant.

Experiment : Initial Deflection Regulation

In this experiment, the states of the plant are given some initial value. The reference model in Figure 1 is a second order system with non-zero initial conditions and high damping (damping ratio $\zeta = 1.0$). The adaptive controller forces the plant output to track the highly damped reference model output. The initial value of S was chosen to be $\{\sigma_1 = 0.5, \sigma_2 = 6.0, L = \bar{L} = 50\}$. Table 1 shows the results of the experiment.

	PI	RMS	ME	ST
open-loop	9.61	0.058	0.139	*
closed-loop without learning	1.82	0.012	0.021	4.35
closed-loop with learning	1.39	0.008	0.019	4.45

Table 1 Experiment results

* The open-loop response did not settle during the simulation time and the ST is set to 24.5 when the PI is calculated.

After learning, the parameter set (optimal) is adjusted to be

$$S_{opt} = \{\sigma_1 = 0.061, \sigma_2 = 8.49, L = 115.5, \bar{L} = 178.4\}$$

Note that the PI has improved by approximately 24% and the percentage changes to the parameters vary from approximately 41% to approximately 257%. The RMS decreases from 0.012 before learning to .008 after learning, for a 33.3% change, the biggest in the three. ST increased by 0.1 for a 2.3% change. This is due to the different weighting in the definition of the PI in equation (5), which assigns the largest weight, 100, to the RMS and smallest weight, 0.1, to the ST. Thus it is not surprising that although PI decreased after learning, some of the individual terms of the PI increased a bit. The sensitivity table in Table 2 shows that $\bar{L}$ is the most sensitive parameter over all and that L is second most sensitive parameter. In general, the parameters related to the proportional gain are more sensitive to the ones related to the integral gain. Columnwise we can see that the RMS and the ST are the most sensitive to $\bar{L}$ and the ME is the most sensitive to L.

	RMS	ME	ST
σ_1	0.03	0.04	0.00
σ_2	0.01	0.08	0.00
L	0.10	0.26	0.01
$\bar{L}$	0.17	0.25	0.05

Table 2 Sensitivity Table

In the experiment, we found that the initial value of S is important for the learning system to work properly. This is especially important for the parameters L and $\bar{L}$. If they are too large, the system becomes unstable. If they are too small, the controller does not provide sufficient input to the plant to affect its behavior. Our assumptions on the performance surface are reasonable when the initial value of S is relatively close, within a factor of one to two, to the optimal set. If S starts far away from S_{opt}, the system may reach a local minimum. This problem will be studied in the future research.

The experimental results of this learning system are quite encouraging. The important achievement is that the learning system found the S_{opt} by itself successfully. The system can of course be improved in many aspects, e.g. the local minimum problem could be addressed and the efficiency of the algorithms could be explored. But it shows that its applications in solving many control problems are promising. For example, this learning system can be used in different operating regions for the adaptive control system, and a decision making mechanism can be built for the on-line tuning of the parameters in the controller based on the experiment results. This will increase the adaptability of the system, or in some sense, add "intelligence" to the system.

6.0 Towards an Intelligent Adaptive Control System

In conventional adaptive control systems, the variation of system dynamics that can be dealt with is quite limited. This is a restrictive limitation because it is desirable, especially in space, for the adaptive control system to be able to deal with wide ranges of environmental changes without human supervision. One way to avoid this limitation of the conventional adaptive control systems is to add some "intelligence" to the system so that it is able to tune its own parameters automatically[1]. To achieve this goal, a real-time knowledge-based system can be used to monitor the adaptive control system.

This real-time knowledge-based system, which has a database, a rulebase, and an inference engine, is built based on the results from the learning discussed above. The database is constructed by running the learning system in different operating regions of the adaptive control system. The database consists of frames with each frame containing information such as the type of disturbance, the reference model, the corresponding optimal parameter set S_{opt}, and the sensitivity table. The rulebase is empty at the beginning. When the inference engine receives this information and together with the performance information (the PI as well as the RMS, the ME, the ST) from the control system, it locates the corresponding S_{opt} and sends it to the control system. Meanwhile, it generates the rules for the on-line turning by reading the look-up table in the same frame. For example, from reading Table 2, the inference engine would generate rules such as:

If RMS is above the threshold, then tune $\bar{L}$.

The rules are fired during the control execution when the performance information indicates that some of the components in the PI, i.e. the RMS, the ME, or the ST, are not tolerable, and an immediate on-line tuning of the parameter set is required. In tuning the parameter, the heuristics of control design can also be incorporated into the rulebase.

This real-time knowledge-based expert system is also implemented in the expert system shell CxPERT. Further on-line testing will be performed as future research.

7.0 Conclusions

The application of learning methodologies in AI provides a method to solve the parameter choosing problem in the adaptive control system of a large space structure. The results of the implementation are encouraging. It is possible that this approach can be applied to other adaptive control systems where the selection of the controller parameters is crucial for the successful implementation. It is also proposed that the results of the learning can be used in automatic on-line tuning system for the intelligent adaptive control system.

Acknowledgment

This work was supported in part by Jet Propulsion Laboratory, Pasadena, CA, under contract No. 957856. Also, the authors would like to thank Kevin Passino and Michael Sartori for their helpful suggestions in preparing this document.

References

[1] Antsaklis, P.J., Passino, K.M. and Wang, S.J., "Autonomous Control Systems: Architecture and Fundamental Issues", Proc. of the American Control Conference, Vol. 1, pp. 602-607, June 1988.

[2] Bazaraa, M.S., and Shetty, C.M., Nonlinear Programming Theory and Algorithms, John Wiley & Sons, New York, 1979.

Z. Gao, M. Peek and P. J. Antsaklis, "Learning for the Adaptive Control of a Large Flexible Structure," Proc. of the 3rd IEEE International Symposium on Intelligent Control, pp. 508-512, Arlington, VA, Aug. 24-26, 1988.

[3] Bayard, D.S., "Averaging Methods for Transient Analysis in Adaptive Control of Large Space Structures", Jet Propulsion Laboratory Engineering Memo, EM 347-88-228, Feb. 4, 1988.

[4] Bayard, D.S., "Closed-Form Bounds on the Adaptive Transient Response of Large Space structures", Jet Propulsion Laboratory Engineering Memo, EM 347-88-235, July 6, 1988.

[5] Charniak, E., and McDermott D., Introduction to Artificial Intelligence, Addison-Wesley Publishing Company, Reading, Massachusetts, 1985.

[6] Cohen, P.R. and Feigenbaum, F.A., The Handbook of Artificial Intelligence, Vol. 3, William Kaufmann, Los Altos, California, 1982.

[7] Fu, K.S., "Learning Control Systems--Review and Outlook", IEEE Transactions on Automatic Control, Vol. AC-15, pp. 210-221, April 1970.

[8] Ih, C.-H.C., Bayard, D.S., Wang, S.J., and Eldred, D.B., "Adaptive Control Experiment with a Large Flexible Structure", to appear in the Proc. of AIAA Guidance, Navigation, and Control Conference, August, 1988.

[9] Mendel, J.M., and Fu, K.S., Adaptive, Learning and Pattern Recognition Systems, Academic Press, New York, 1970.

[10] Michalski, R.S., Carbonell, J.G., and Mitchel, T.M., Machine Learning - An Artificial Intelligence Approach, Tioga, Palo Alto, 1983.

[11] Sklansky, J., "Learning Systems for Automatic Control", IEEE Transactions on Automatic Control, Vol. AC-11, pp. 6-19, Jan. 1966.

[12] Tsypkin, YA.Z., Adaptation and Learning In Automatic Systems, Academic Press, New York, 1971.

[13] Tsypkin, YA.Z., "Self-Learning-What Is It ?", IEEE Transactions on Automatic Control, Vol. AC-13, pp. 608-612, Dec. 1968.