JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 24, No. 5, September–October 2001

Limit Cycle Oscillation Prediction Using Artificial Neural Networks

Charles M. Denegri, Jr.* and Michael R. Johnson†
Air Force SEEK EAGLE Office, Eglin Air Force Base, Florida 32542-6865

A static artificial neural network in the form of a multilayer perceptron is investigated to determine its ability to predict linear and nonlinear flutter response characteristics. The network is developed and trained using linear flutter analysis and flight-test results from a fighter test. Eleven external store carriage configurations are used as training data, and three configurations are used as test cases. The network was successful in predicting the aeroelastic oscillation frequency and amplitude responses over a range of Mach numbers for two of the test cases. Predictions for the third test case were not as good. Several network sizes were investigated, and it was found that small networks tended to overgeneralize the training data and are not capable of accurate prediction beyond the sample space. Conversely, networks that were too large, or trained to error levels that were extreme, tended to memorize the training data, and are also unable to produce adequate predictions beyond the sample space. The results of this study indicate that relatively simple networks using small training sets can be used to predict both linear and nonlinear flutter response characteristics.

Introduction

LIMIT CYCLE oscillations (LCO) have been a recurring problem on certain fighter aircraft and are generally encountered on external store configurations that are theoretically predicted to be flutter sensitive. These sensitivities are quite evident during flight and are often the subject of extensive examination during flutter flight tests of aircraft that exhibit this behavior. Reference 1 provides a detailed description of the LCO phenomenon and its relationship to classical flutter. An excellent overview of LCO of fighter aircraft carrying external stores and its sensitivity to the store carriage configuration and mass properties is given in Ref. 2. These articles describe LCO as a phenomenon characterized by sustained periodic oscillations that neither increase nor decrease in amplitude over time for a given flight condition. These articles also describe the problems associated with this phenomenon and the elusiveness of predicting its occurrence theoretically.

LCO arises from the nonlinear interaction of the structural and aerodynamic forces acting on the affected aircraft component. Several approaches exist for predicting the occurrence of LCO for fighter aircraft. The most practical approaches[3–5] are empirical in nature and are based on the assumption that LCO is a variety of classical flutter, that is, once oscillations initiate they catastrophically diverge. LCO differs from classical flutter in its tendency toward limited amplitude oscillations rather than diverging oscillations. The assumption that LCO is a form of classical flutter is substantiated by the fact that the occurrence of LCO is usually associated with flutter-sensitive aircraft/store configurations. Linear flutter analysis theory forms the foundation upon which current LCO prediction methods are based. A brief description of typical linear flutter analysis and LCO analysis follows.

Linear flutter analyses are typically accomplished in the frequency domain and involve obtaining the natural vibration frequencies and modes, calculating the generalized aerodynamic forces, and then solving for the modal damping and frequency variations with velocity. This analysis assumes the response of the structure to be simple harmonic motion and is therefore limited to indicating the stability of particular modes with respect to flight velocity. The main shortcoming of these methods with regard to LCO is that no indication of the oscillation amplitude is available. This parameter is of primary importance in the certification process because configurations that exhibit high-amplitude, neutrally stable oscillations are typically avoided while those exhibiting low-amplitude oscillations are typically deemed suitable for use. Even though linear flutter analyses are not capable of directly predicting LCO, these analyses have been shown to adequately identify LCO-sensitive store configurations and the oscillation frequency of the instability. Furthermore, prior studies[6,7] have shown that the modal composition of the LCO mechanism can be indicative of the general nature of the LCO sensitivities.

In contrast to linear flutter analyses, LCO analysis is typically accomplished in the time domain and involves perturbing the structure, computing the resulting aerodynamic forces, and then solving for the resultant structural deformation caused by these forces. The aerodynamic forces are then recomputed, and the process steps forward in time. In this manner a time history of the structural response is predicted from which the damping and amplitude characteristics of a particular configuration can be examined.

One of the most restrictive problems with performing LCO analyses is the tremendous volume of analysis cases that must be examined in order to provide certification for a given store carriage configuration and its permutations. The large number of possible store carriage configurations (as a result of external store downloading) increases the likelihood of encountering LCO in the flight envelope for some store carriage permutations. From a structural dynamics perspective different external store configurations essentially alter the mass and inertia characteristics of the wing structure. If different weapon carriage pylons are involved, then the aeroelastic stiffness characteristics are altered as well. The presence of stores and pylons can also significantly affect the aerodynamics. From these considerations it is apparent that because multiple stores can be carried and downloaded on a typical fighter aircraft, then numerous aeroelastic systems exist for each store carriage configuration. Thus, enormous quantities of flutter or LCO analyses must be performed for a single aircraft weapon configuration.

Because typical flutter and LCO analyses only give an indication of the potential in-flight behavior, flight testing of the most critical configurations is accomplished in order to verify the analyses and to determine the true LCO characteristics. Flight testing continues to become more and more expensive, whereas the number and combinations of store configurations are continually increasing. Hence, it is important to find methods that accurately determine the flutter and LCO characteristics of potentially dangerous configurations without the need for testing. It is equally desirable to identify those

Presented at the CEAS/AIAA/ICASE/NASA Langley International Forum on Aeroelasticity and Structural Dynamics, Williamsburg, VA, 22–25 June 1999; received 13 August 1999; revision received 20 November 2000; accepted for publication 13 December 2000.

*Lead Flutter Engineer, Engineering Division, 205 West D Avenue, Suite 348. Senior Member AIAA.

†Lead Ballistics Engineer, Analysis Division, 205 West D Avenue, Suite 348. Senior Member AIAA.

configurations whose response characteristics are acceptable, and thereby eliminate the need for unnecessary testing. Although some success in theoretically predicting LCO has been achieved, the primary shortcoming of these methods is that they have not yet been shown to be practical for applications that require a large number of analyses (such as weapon certification efforts on fighter aircraft). The present work examines the feasibility of mapping the results of previous flight tests into an intelligent system that can infer the results of slightly differing configurations before any test missions are flown.

Fundamental Assumptions

In subsequent sections the feasibility of using an artificial neural network (ANN) for predicting LCO on a fighter aircraft with external stores is evaluated. In essence, an ANN is constructed based upon empirical inputs from linear flutter analysis and LCO flight tests. The network is then trained to provide an LCO prediction capability for new configurations using linear flutter analysis data as inputs. The goal of this analysis approach is to develop a practical nonlinear LCO prediction capability, which provides greater insight into the expected test results, thus reducing the need for expensive flight testing. It must be emphasized that no attempt is made in the present work to explain the physics of the LCO phenomenon. The method presented only attempts to predict the occurrence of LCO based on historical flight test results by mapping inputs to known flight-test results and using the interpolative ability of the network to "infer" the response of similar, but untested configurations. This analysis approach requires a sufficient historical database of the LCO characteristics of the particular aircraft for which predictions are sought. The fundamental assumptions and supporting rationale for this analysis approach are discussed in the following.

The first fundamental assumption is that LCO is considered to be a variation of classical divergent flutter whereby the inertial, elastic, and aerodynamic forces combine to initiate unstable oscillations. This assumption is supported by the observation that linear flutter analyses are capable of determining LCO-sensitive store carriage configurations, and flight testing shows strong similarities between LCO and classical flutter responses.[6,7] The only difference in the responses is that for LCO the oscillations stabilize at a limited amplitude rather than diverging.

The second fundamental assumption is that LCO is considered to differ from classical flutter in that the nonlinearities present in the system can only serve to limit the amplitude of the resulting oscillations and do not change the flutter speed of the aircraft configuration.[8] Such LCO are termed benign and can arise from an increase in structural stiffness or damping with increasing amplitude of oscillation[9] or a favorable effect of some fluid nonlinearity. These benign LCO are distinct from other LCO that can occur at flow conditions below the classical linear flutter speed, for example, caused by freeplay in a control surface.[10] This assumption is supported by observations during flight tests of LCO-prone fighter aircraft. These observations show that the flutter speed is relatively independent of the strength of the disturbances used to excite the aeroelastic system, which implies that the nonlinearities are benign.

The third fundamental assumption is that the physical nonlinearities (whether structural or aerodynamic), which cause LCO in a new configuration for which a prediction is sought, must be adequately represented in the historical data. The nonlinearities of the new configuration are considered to be essentially the same as for the flight-test configurations used in training the ANN. Furthermore, it is assumed that these nonlinearities do not vary from one configuration to another, which implies that, for a given aircraft, the nonlinearity is not related to a specific store configuration. That is, it is assumed that the store configuration can change the flutter speed or LCO onset speed, but not the nonlinearity per se. This implies that the nonlinearity is inherent in the wing structure or aerodynamics.

The last fundamental assumption is that components of the linear flutter analyses are considered to be adequate for predicting the general nature of the aeroelastic instabilities and to predict, qualitatively at least, the onset of LCO. There is evidence that subtleties in the linear flutter analysis results correlate to distinguishing characteristics of the LCO behavior.[6,7] For example, the speeds and frequencies of the critical damping crossings and the modal composition of the linear flutter analysis mechanism are generally accepted as being related to the LCO behavior in flight.

These assumptions are not unreasonable, but it is recognized that they cannot be universally true. However, with the aforementioned assumptions in mind, the results presented here offer insight to those attempting to build mathematical models from first physical principles.

Analysis Approach

The concept of the ANN is an attempt to simulate one popular model of the memory structure of the human brain. The ANN is designed to reproduce the brain's behavior in terms of learning and adaptation. The desirable characteristics of the ANN lie in its ability to identify and model highly nonlinear systems. ANNs have been shown to exhibit a potential for highly effective interpolation[11] within a problem space and can be used as a tool for the prediction of nonlinear states beyond the problem space bounds. It is the ability to predict system behavior that makes ANNs attractive for the prediction of LCO.

A neural network is an interconnection of nodes and weights that map an input vector to an output vector. Nodes, sometimes referred to as perceptrons, are traditionally comprised of two parts: a summing connection where values from incoming signals are added and a function that transforms the summed signal into another value. Strictly speaking, a perceptron is a specific combination of a summing node and a hard-limiting transfer function that maps the input to one of two output states.[12] The term has come to be used interchangeably with *node*, however. A node is illustrated in Fig. 1. The transfer functions can be any mathematical function that transforms, or maps, a given input to a given output, but the power in the ANN lies in the use of nonlinear functions, giving the network the ability to map linear combinations of input vectors into nonlinear results. Typical nonlinear functions include exponential or logarithmic sigmoids and hyperbolic tangents. Output layers are more likely to be linear, either limited by upper or lower bounds or both, or unlimited. Figure 2 shows a small, two-input, single-output ANN. Two nonlinear node functions are shown in Fig. 3.

Neural networks contain layers of nodes between the input and output, referred to as hidden layers because they are invisible to the user and autonomously perform their functions. A network can contain an unlimited number of hidden layers, but typically no more than two are required.[11] A static neural network is one in which the input vector is fed forward only. A dynamic neural network contains feedback from either the output back to the input or between layers hidden within the network. The advantage of a static network is the relative simplicity by which it can be adapted. Dynamic networks, on the other hand, must make use of adaptation processes that, in effect, unfold the network into its static equivalent.

Static networks are an obvious choice for modeling finite impulse response systems. They also work very well for mapping from one space to another (hence the use of a static ANN in this work). If

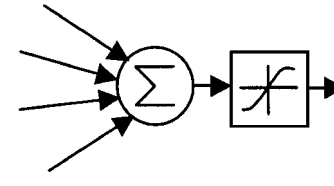

Fig. 1 Single network node.

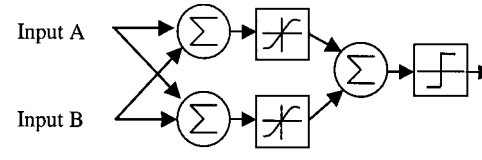

Fig. 2 Simple two-node network.

Fig. 3 Network nonlinear functions.

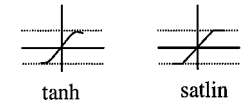

the system being modeled has a transfer function that has infinite impulse response (IIR) characteristics, a dynamic network may be more applicable. Because of the IIR nature of the dynamic network, it is typically used when modeling a temporal process. It can also be used, however, to model any process whose input and desired response are sequential in nature.

The process of backpropagation is used to train, or adapt, a neural network. The network "learns" by repeated exposure to the desired response. The network is given an input vector, and a response is measured and compared to the desired response. The difference between the two is an error signal, which is used to readjust the network values. Backpropagation is the process by which the error signal is propagated back through the various layers of the network. A detailed description of backpropagation is beyond the scope of this work, but refer to Ref. 11 for details. Dynamic networks require more complicated processes for adaptation. Backpropagation through time is used to effectively unfold the network into equivalent static layers that would be encountered at each time step in the input vector.

A static neural network in the form of a multilayer perceptron (MLP) was chosen for this study because of its ease of design and evaluation, as well as its simplicity. A static network is designed to simply feed forward input sets one at a time and provide a prediction. By its nature, there is no inherent knowledge of past inputs or outputs. The problem then is reduced to one of functional representation. It is from this perspective that the ANN is employed for LCO prediction.

The ANN is trained using both flight-test data and linear flutter analysis data. The flight-test data from a wide variety of external store configurations are used and presented in Table 1. The external store configurations are presented in Table 2, and the flutter analysis data are presented in Table 3. These configurations exhibit characteristics that are representative of the broad spectrum of flutter and LCO responses encountered by fighter aircraft with external stores. Reference 6 categorizes these characteristics and shows a correlation between the modes comprising the predicted linear flutter mechanism and the flight response characteristics. Based on these results, the dynamic characteristics of the aircraft store configuration are represented as inputs to the ANN by the quantized free vibration mode shapes and frequencies that comprise the linear flutter analysis predicted mechanism. This is done in an effort to ensure the applicability of the method to store configurations that have not been flight tested. Finally, the aerodynamic characteristics are represented to the ANN by the linear analysis flutter speed and frequency and by a quantized representation of the wing-tip store configuration. In this manner all of the fundamental inertial, elastic, and aerodynamic characteristics of the aeroelastic system are represented. Known LCO response level and response frequency are used for output training. The networks are trained using constant amplitude response data. It follows then, that they seek to predict a constant amplitude response. Zero-amplitude response in either the flight-test training data or the network output implies either no response or damped response. Therefore, it is possible for an oscillatory frequency to be associated with zero-amplitude response.

Table 1 Flight-test response amplitude[a]

	Mach number										
Configuration	0.80	0.85	0.90	0.91	0.92	0.93	0.94	0.95	0.98	Frequency, Hz	Category
1	0.0	0.5	1.5	——	——	——	——	2.5		6.60	Typical LCO
2	0.0	0.0	0.0	——	——	2.5	4.0			9.40	Flutter
3	0.0	0.0	0.0	——	——	——	4.5			9.40	Flutter
4	0.0	1.0	2.0	3.0						6.80	Typical LCO
5	0.0	0.0	1.5	——	——	——	——	0.0		7.00	Nontypical LCO
6	0.0	0.0	0.5	——	——	——	——	1.0	1.5	6.90	Typical LCO
7	0.0	0.0	0.0	——	——	——	——	0.5		9.20	Flutter
8[b]	0.0	0.0	1.0	——	2.0					9.50	Flutter
12	0.0	1.0	2.0	——	——	——	——	2.5		6.80	Typical LCO
13[b]	0.0	1.0	3.0	4.0						7.80	Typical LCO
15	0.0	0.5	1.0	——	——	2.0	——	0.0		8.10	Nontypical LCO
16	0.0	0.0	0.0	——	——	——	——	0.5		7.00	Typical LCO
17[b]	0.0	0.0	0.5	0.0						8.20	Nontypical LCO
18	0.0	1.0	2.5	3.0						8.10	Typical LCO

[a]Amplitude in units of gravitational acceleration g. All amplitudes are for constant oscillatory response except for configurations 2, 3, 7, and 8, which showed divergent response at the endpoints.
[b]Network test configuration.
—— indicate no response data explicitly measured.
Blanks indicate no test data acquired.

Table 2 External store configurations

	Station 1 Wing-tip missile		Station 2 Underwing missile		Station 3 Weapon		Station 4 Fuel tank
Configuration	Suspension equipment	Store	Suspension equipment	Store	Suspension equipment	Store	Fuel state
1	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 3	$\frac{1}{4}$-full
2	Launcher A	None	Launcher A	None	Launcher C	Missile 3	No tank
3	Launcher A	None	None	None	Launcher C	Missile 3	Empty
4	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 1	$\frac{1}{4}$-full
5	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 1	Empty
6	Launcher A	None	Launcher A	Missile 1	Launcher C	None	$\frac{1}{2}$-full
7	Launcher A	None	Launcher A	None	Launcher C	Missile 3	Empty
8[a]	Launcher A	None	None	None	Launcher C	Missile 3	No tank
12	Launcher B	None	Launcher A	Missile 1	Launcher C	Missile 3	$\frac{1}{4}$-full
13[a]	Launcher A	None	Launcher A	Missile 2	Launcher C	Missile 3	Empty
15	Launcher A	None	Launcher A	Missile 2	Launcher C	Missile 1	Empty
16	Launcher B	None	Launcher A	Missile 1	Launcher C	None	$\frac{1}{2}$-full
17[a]	Launcher B	None	Launcher A	Missile 2	Launcher C	Missile 1	Empty
18	Launcher B	None	Launcher A	Missile 2	Launcher C	Missile 3	Empty

[a]Network test configuration.

Table 3 Linear flutter analysis, 0.95 Mach, sea-level aerodynamics

	Flutter speed, KCAS						
	Damping		Flutter	Flutter	Natural	Coupled	Natural
Configuration	0%	1%	frequency, Hz	mode	frequency, Hz	mode	frequency, Hz
1	467	537	6.89	FWB[a]	7.03	AWB[b]	6.34
2	726	776	9.09	1WB[c]	8.88	1WT[d]	9.76
3	635	689	9.21	1WB	9.02	1WT	9.89
4	473	533	6.92	FWB	7.07	AWB	6.45
5	876	924	7.51	1WB	7.80	1WT	6.93
6	509	566	6.90	FWB	7.09	AWB	6.49
7	658	702	8.91	1WB	8.68	1WT	9.69
8[e]	726	790	9.35	1WB	9.19	1WT	9.96
12	435	516	6.95	FWB	7.07	AWB	6.40
13[e]	327	519	8.09	FWB	8.14	AWB	7.89
15	449	616	8.15	AWB	8.13	FWB	8.25
16	493	554	6.97	FWB	7.14	AWB	6.55
17[e]	455	653	8.25	1WT	8.31	1WB	8.28
18	291	538	8.19	FWB	8.23	AWB	7.98

[a]Forward wing bending. [b]Aft wing bending. [c]First wing bending. [d]First wing torsion. [e]Network test configuration.

The frequencies associated with zero-amplitude responses are the dominant frequencies of the damped responses.

Some network inputs are quantized values of descriptive information. The network requires numerical quantities as input. Thus, all of the descriptive data is represented as integers. The method used for quantization is rather arbitrary. For this work the descriptors were simply ordered with integer values ranging from 0 to $n-1$ (where n is the total number of descriptors being quantized). For example, Table 2 shows that there are two different launchers that can be carried on the wing tip (station 1). These are launcher A or launcher B. These descriptors were quantized with the value 0 corresponding to launcher A and 1 corresponding to launcher B. A similar process was used for all other descriptive variables. The fuel state of the tank on station 4 was assigned values from 0 to 3. The four mode shapes were quantized with 0 for forward wing bending, 1 for first wing torsion, 2 for aft wing bending, and 3 for first wing bending.

After the ANN has been trained, it is tested using selected LCO cases from Ref. 7. The ANN gives, as output, the LCO response amplitude and frequency as a function of Mach number. The analysis results and the correlation between the size of the training set and the convergence of the method are discussed.

LCO Characteristics

Reference 7 describes three categories of response behavior seen on fighter aircraft: flutter, typical LCO, and nontypical LCO. Classical flutter behavior is characterized by divergent wing oscillations. (In practice, the sudden onset of high-amplitude wing oscillations that show no tendency toward a limited amplitude is interpreted as flutter behavior.) Typical LCO is characterized by the gradual onset of sustained limited-amplitude wing oscillations, where the oscillation amplitude progressively increases with increasing Mach number and dynamic pressure. Nontypical LCO is characterized by the gradual onset of sustained limited-amplitude wing oscillations, where the oscillation amplitude does not progressively increase with increasing Mach number and oscillations may be present only in a limited portion of the flight envelope.

The linear flutter analysis results for each configuration are presented in Table 3. These analyses are *not* matched analyses but merely worst-case "screening" analyses. In this manner all analyses are performed using sea-level density and 0.95 Mach doublet-lattice method[13] aerodynamic influence coefficients. The free vibration analyses are performed for a half-airplane model using a matrix iteration method and the first 16 antisymmetric flexible modes are retained for the flutter solution. The flutter equations are solved using the Laguerre iteration method,[14] which is a variation of the classical k-method of flutter determinant solution. The structural and aerodynamic models used for these analyses are shown in Fig. 4.

For these flutter analysis results a critical point is considered to be the velocity at which a modal stability curve crosses from stable (negative damping required to produce neutral stability) to unstable (positive) damping. The analytical flutter speed is the critical point associated with the known aeroelastically sensitive mode for the particular configuration. For comparison purposes the analytical flutter speed is considered to be directly comparable to the lowest airspeed at which self-sustained oscillations are encountered in flight. These oscillations could be either LCO or flutter. The slope of the modal damping curve indicates the velocity sensitivity of a mode. The velocities at the 0 and 1% damping levels determine this slope. Steep slopes indicate rapid decreases in stabilizing damping with increased velocity.

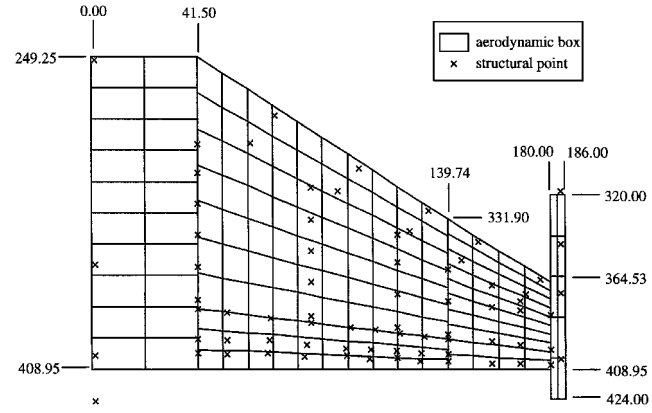

Fig. 4 Flutter analysis model composed of doublet-lattice aerodynamics and lumped mass structure (all dimensions in inches).

The flight test results are presented in Table 1, where it is seen that each configuration exhibits one of the three previously discussed categories of response behavior (flutter, typical LCO, or nontypical LCO). The flutter analysis results (Table 3) show that for each of the three response categories a common set of modes are present in the linear flutter mechanism. Reference 6 showed a distinct correlation between these modes and the flutter or LCO behavior. Essentially, it was shown that different linear analysis flutter mechanisms correlated to different aeroelastic responses in flight. Thus, the mode shapes and the frequencies of the critical modes (Table 3 and Fig. 5) are used as primary inputs to the ANN. The flight-test results (Table 1) are for level flight at 5000-ft (1525-m) pressure altitude. The flight tests were generally conducted in 0.05 Mach increments beginning at the lower Mach number. Smaller increments were used when large response amplitudes were encountered or expected. A test point maneuver was terminated when the response amplitude either exceeded predetermined termination criteria, or the response amplitude increased at such a rate as to rapidly approach the predetermined termination criteria. So, the values presented may not be absolute maximum response levels, but merely the highest response level measured before terminating the test point. Details on the flight test procedures can be found in Ref. 7.

Network Design

For the present work an ANN is designed in the form of a MLP as shown in Fig. 6. This consists of an input layer, two hidden layers, and an output layer. The inputs are directed to summing nodes through the input weights, whose sum is fed into nonlinearities in the form of hyperbolic tangents (tanh). The outputs of the tanh nodes are then weighted again, summed, and fed into another tanh layer. The process is repeated again, this time through saturated linear (satlin) output nodes. Both the tanh and satlin functions have the property of limiting outputs to the range [−1, 1]. These functions are shown in Fig. 3.

The size of the weight matrices and the number of nodes necessary to successfully represent the underlying nonlinear function are not easily determined. The network must be large enough to fully absorb all training data and allow reasonable interpolation and extrapolation, yet be small enough to be manageable. On the other hand, networks that are too large or trained to very small mean squared-error values tend to memorize the training data and are of little use for generalizing a function. Following general design rules of thumb found in Ref. 11, many network sizes were considered.

Supervised learning, that is, presenting the desired output to the network in order to generate an error signal, which is then fed back through the network by backpropagation, is used to train the network. The networks were batch trained using the Levenberg–Marquardt[15] algorithm, which uses the Jacobian matrix containing the first derivatives of the network errors with respect to the weights and biases. The weights are updated using the equation

$$W_{k+1} = W_k - (J^T J + \mu I)^{-1} J^T e$$

where W are the weights, J is the Jacobian matrix, I is the identity matrix, e is the error vector, and μ is a weighting factor (less than 1) that decreases proportionally to e as the solution approaches a minimum. Training continues until the average mean squared error falls below a preset value.

All networks considered used the data shown in Tables 1–3 and Figs. 7–9. Thirteen inputs were given for each output frequency and amplitude vs Mach combination (Fig. 6). Inputs are quantized representations of the aircraft configuration, flight Mach number, and linear flutter analysis results for the configuration. All input data were normalized to span the range [−1, 1], consequently the

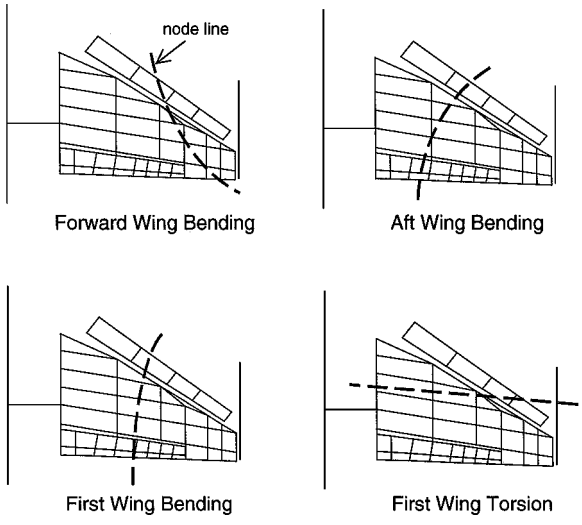

Fig. 5 Linear flutter analysis mode shapes (see Table 3).

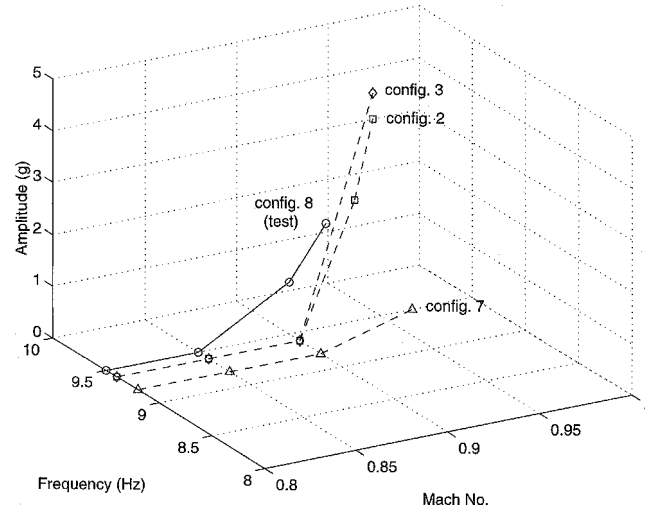

Fig. 7 Flight-test response characteristics, flutter cases.

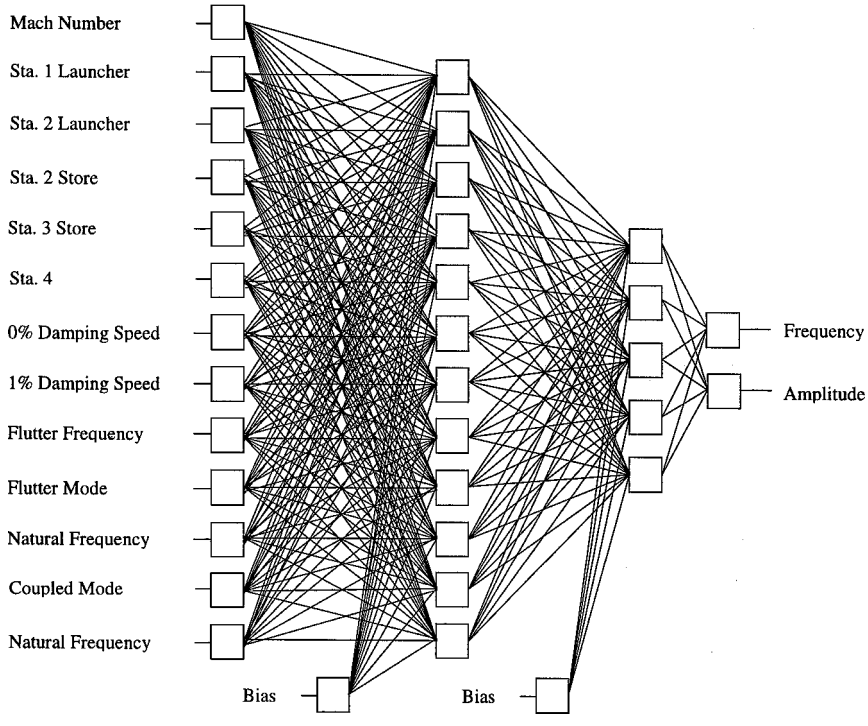

Fig. 6 12 × 5 network diagram.

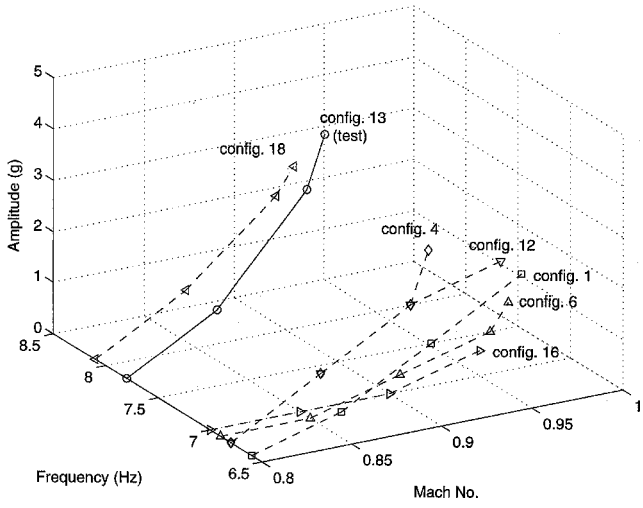

Fig. 8 Flight-test response characteristics, typical LCO cases.

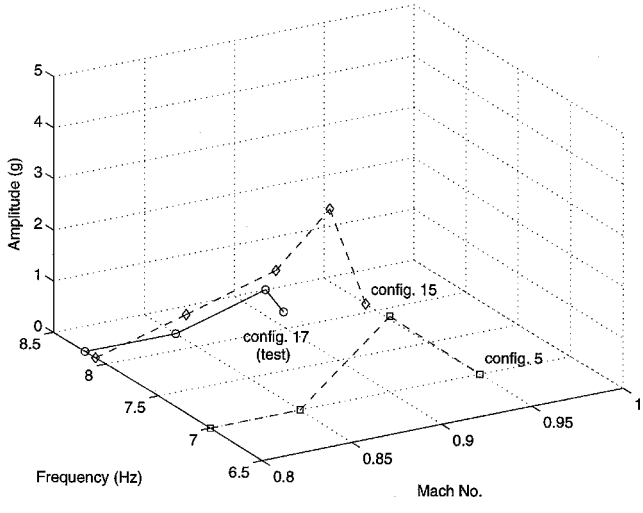

Fig. 9 Flight-test characteristics, nontypical LCO cases.

outputs fall within that range as well and have to be scaled back to workable values. Over all, there are 14 flight-test/linear model combinations available for training and testing the network. Three of these combinations were held aside for testing and evaluating the effectiveness of each network. These were flight-test configurations 8, 13, and 17 (Figs. 10–12) and represented a flutter case, a typical LCO case, and a nontypical LCO case. Comparison of the training and testing curves for the flight-test data are shown in Figs. 7–9. These figures show that the amplitude, frequency, and Mach-number characteristics of the test cases are sufficiently different from the training data to be a difficult predictive test for the network. The training and testing data are shown grouped according to their response behavior category. However, the network is trained using all of the data except the network test configurations. It is left up to the network to determine the amplitude vs Mach-number characteristics that then define the aeroelastic response category of the test configuration results.

The network is trained using the squared-error fed back as described earlier. Initially the minimum value used to indicate a successfully trained network was set at $e = 10^{-6}$. A large network with 30 nodes in the first hidden layer and 12 nodes in the second (30×12) was initially trained. This network fit the training data very well but did not predict the test data set, indicating that the network was too large and overtrained. In effect, it had memorized the training data. Upon further consideration, the decision was made to set the minimum squared-error value at 10^{-2}. This more closely reflects the data characteristics, in that both the desired oscillation frequency and amplitude are accurate to one significant digit. Consequently, the squared error is on the order of 10^{-2}. Subsequent trial networks were trained to this value.

Several network sizes were investigated and are discussed without proof. The smallest network considered had a single layer of five nodes. The largest was a 50-by-50 double hidden layer MLP. The small networks tended to overgeneralize the data. That is, the output tended to be "too smooth" and did not follow the trend of the training data very well at all. The large networks tended to memorize the data and did not give reasonable extrapolation outside the solution space. Each network was trained several times in order to increase the probability of truly finding the global minimum in the error space.

Different output node types were evaluated as well. Three were considered: the tanh function; the saturated linear (satlin) node; and a linear node (output range $[-\infty, \infty]$). The tanh function performed

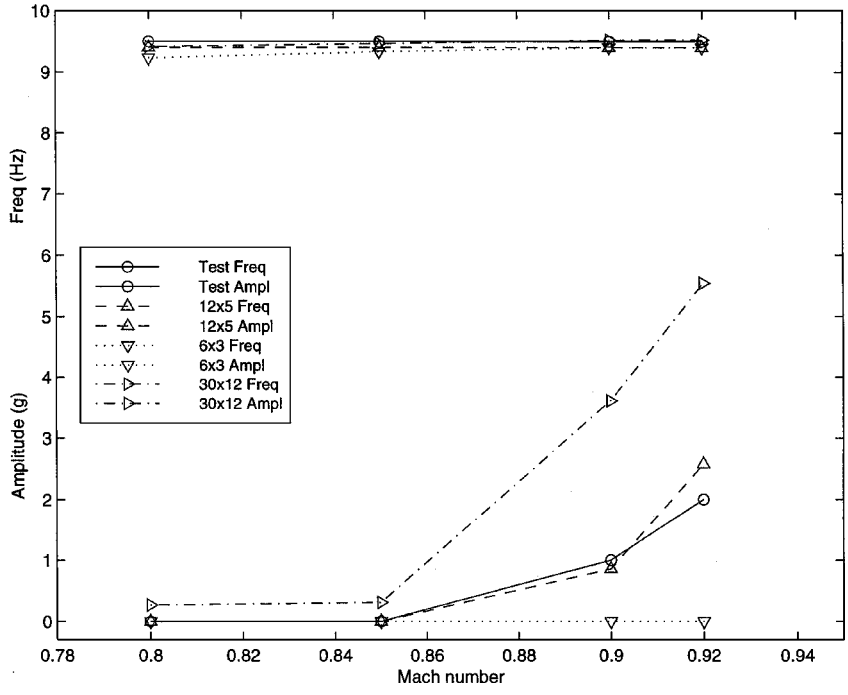

Fig. 10 Flutter test case.

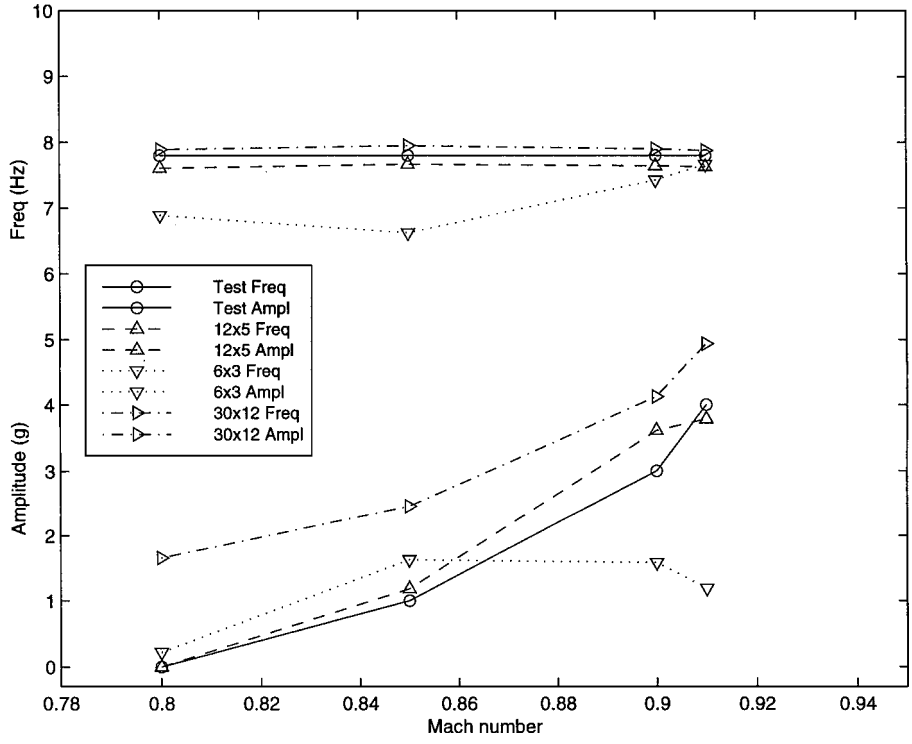

Fig. 11 Typical LCO test case.

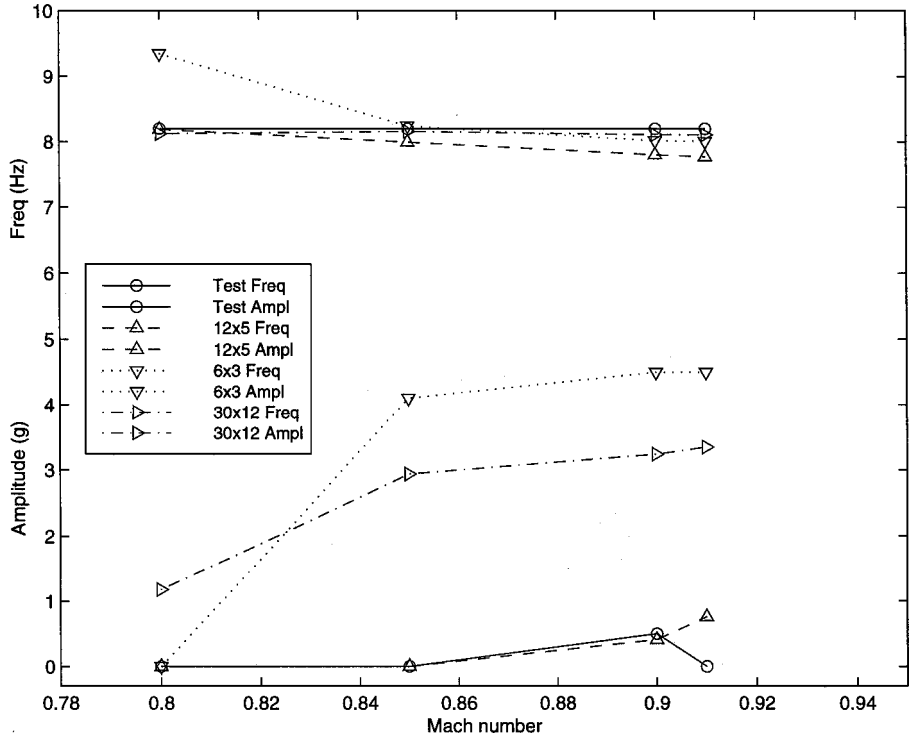

Fig. 12 Nontypical LCO test case.

similarly to the satlin. Because the output data in this study are assumed linear, the design tended toward the satlin function. This function gives true linear characteristics, while bounding the solution to the defined solution space.

The final design scheme selected for the static MLP contains 12 tanh nodes in the first hidden layer and five in the second, with two satlin nodes as outputs (12×5) and is shown in Fig. 6. This configuration gave good generalization over the solution space, while learning the training data well. The outputs are the combined oscillation frequency in hertz and oscillation amplitude in units of gravitational acceleration g at the forward end of the wingtip launcher. All input data are normalized to $[-1, 1]$. The output is also contained in $[-1, 1]$ as a result of the saturated linear output nodes. All output data are then processed back to usable values by reversing the normalization process.

Results

Network inputs for the test cases consisted of the Mach number of the desired flight condition, the store carriage configuration, and the linear flutter analysis results (flutter speed and frequency, the modal composition of the flutter mechanism, and the free vibration frequencies of those modes). The output from the network was in the

form of an oscillation response amplitude and frequency. The three-dimensional flight-test data for the network test cases (Figs. 7–9) is shown in Figs. 10–12 as two two-dimensional curves of frequency and amplitude vs Mach number. From a two-dimensional perspective the test data in Figs. 10–12 appear to be simple polynomials (or straight lines). However, the network outputs must represent the flight-test data as functionally dependent on the linear flutter analysis results and aircraft configuration data as inputs. Thus, this process is not simply curve fitting in two-dimensional space, but rather a multidimensional mapping process with combinations of two-dimensional curves as output.

Several Mach numbers were examined for each set of input conditions, and responses were noted. The classification of the predicted response is described as follows. The network indicated a flutter condition if the output oscillation amplitude increased dramatically as Mach number increased. The network indicated a typical LCO condition if the amplitude increased progressively with no sudden high-level responses. The network indicated a nontypical LCO condition if the amplitude rose to a level, then began to decrease.

For clarity, it is necessary to elaborate here on the rationale for interpreting the ANN output in this manner. The flight-test data used to train the neural network consist of two primary components: the oscillation amplitude and its frequency. Because LCO exhibits a constant amplitude response for a given flight condition, the response characteristics are typically noted in terms of the measured peak sinusoidal amplitude and the frequency of the oscillation. Cases with either damped or random oscillations, that is, no constant sinusoidal response, are denoted in the flight-test database as having a zero-amplitude response, and the frequency associated with this response is the dominant frequency determined from spectral analysis of the response signal. For cases where divergent oscillations are encountered, the flight-test point is terminated as soon as it is apparent that the oscillations are unbounded, and the maximum measured response amplitude is then typically recorded in the database. Because the flight-test points that encountered diverging oscillations were not allowed to diverge fully, the data used in training the network are the actual measured response data, which is not necessarily the maximum response that would have been encountered had the oscillations been allowed to continue. For the LCO cases the maximum constant amplitude response levels that were measured are used to train the network.

Outputs from three of the evaluated networks are shown in Figs. 10–12. The small 6×3 network yielded inadequate results for all three test cases. For the flutter case (Fig. 10) the amplitude computed by this network remained zero for all input Mach numbers. The network shows good amplitude correlation between 0.80 and 0.85 Mach for the typical LCO case (Fig. 11) but fails to track the subsequent amplitude increase. As Mach number is increased, the network diverges from the measured response by showing the amplitude to decrease. Finally, the amplitudes computed for the nontypical LCO case (Fig. 12) were extremely high for all Mach numbers except 0.80 and clearly represent an unsuccessful extrapolation by the network. The small network showed good correlation to the measured oscillation frequency for the flutter case only.

Logically, it was thought that a larger network would perform better than a smaller one in part because of its inherent ability to store more data and cover a larger portion of the solution space. Unfortunately, the observed results were contrary to this. In all three test cases the 30×12 network tracked the oscillation frequency well, but failed to adequately track the response amplitudes. For the flutter case (Fig. 10) it is seen that the computed amplitude has the desired characteristics, that is, the amplitude increases dramatically with Mach number. However, this network tended to yield significantly higher response amplitudes than were measured in flight. The ability of this network to track the shape and trend of the amplitude response is also seen for the typical LCO case (Fig. 11). However, for this case the amplitude is again consistently higher than the measured response. The computed amplitude for the nontypical LCO case (Fig. 12) is excessively high for all Mach numbers.

The 12×5 network showed the best overall agreement with the test data for both oscillation frequency and amplitude. For the flutter test case (Fig. 10) the computed frequency and oscillation amplitude show very good correlation to the measured data. The network amplitude prediction showed a dramatic increase at 0.92 Mach consistent with the flight-measured response.

The typical LCO data available for training were more diverse than for the other cases. This allowed the network's solution space to be larger, but increased the potential for errors in the interpolation between points. For the typical LCO test case (Fig. 11) the network amplitude tracks the measured response very well except for the last point at 0.91 Mach. There the network output shows a trend of leveling off while the measured response amplitude continues to increase. As was observed in the flutter case, it is seen that the frequency of oscillation was predicted nearly exactly.

The nontypical LCO case was the most revealing of the three test cases because it exposed several limitations of the static MLP. As shown in Fig. 12, frequency was not tracked as well as for the preceding cases, indicating that there was not enough information about frequency available to the network for it to form a functional representation. The network amplitude tracked the flight response well up to 0.91 Mach. There, the network indicated a slight increase in amplitude while the flight data showed a decrease. This can be explained by the fact that in the nontypical LCO training cases the amplitude is zero, or near zero, and grows only slightly. Further, the nontypical LCO test case falls outside the sample space of the training data (Fig. 9). This forces the network to attempt to extrapolate beyond its "experience." This is, of course, one of the benefits of using ANNs, but in this case the network was evidently pushed beyond its capability.

Conclusions

From this feasibility study it is concluded that the static ANN was very successful considering the small data set used for training and the limitations of the static network itself. It was shown that both the flutter and typical LCO cases were reasonably predicted with this network. The limitation of the network is shown in the nontypical LCO case. The limited size of the data used for training was such that virtually any prediction outside the space of those two cases constituted an extrapolation into unknown regions of the solution space. As such, the network had little experience to draw from in forming reasonable outputs.

Evaluating the various networks that were considered for this work shows the strengths and weaknesses of their architectures. It was seen that small networks overgeneralize the training data and cannot be used for accurate prediction beyond the sample space. This was demonstrated in the 6×3 network. Conversely, networks that are too large, or trained to error levels that are extreme, tended to memorize the training data, and therefore are also inadequate for prediction beyond the sample space. This was shown to be the case in the 30×12 network trained to $e = 10^{-6}$. The most reasonable results were seen for the moderately sized 12×5 network trained to $e = 10^{-2}$.

Giving consideration to the limited data set used for network training, the results presented offer positive evidence of the feasibility of using an ANN for predicting LCO of fighter aircraft. It is shown that, based on the available training data, flutter behavior was adequately predicted, both in amplitude and in frequency. Typical LCO was also predicted adequately, again by observing amplitude levels. The network had difficulty predicting the (decreasing) change in oscillation amplitude that would indicate nontypical LCO. The training data available for nontypical LCO were limited to two flight-test/linear modeling output data sets, and the test set inputs were outside the training data space. Because prediction of that case constitutes an extrapolation outside the solution space, it is reasonable to expect the errors seen in this study.

The use of the artificial neural network analysis approach provides a significant improvement over classical linear flutter analyses for predicting LCO characteristics. The strengths of linear flutter analyses are that they adequately identify store configurations which are flutter and LCO sensitive and they give a good indication of the instability frequency. However, linear flutter analyses do not accurately identify the instability onset speed and provide no indication

of the response amplitude. By combining the strengths of the classical flutter analyses with historical flutter flight-test results using an ANN, one can obtain adequate LCO predictions in a simple, practical, and timely manner.

References

[1]Bunton, R. W., and Denegri, C. M., Jr., "Limit Cycle Oscillation Characteristics of Fighter Aircraft," *Journal of Aircraft*, Vol. 37, No. 5, 2000, pp. 916–918.

[2]Norton, W. J., "Limit Cycle Oscillation and Flight Flutter Testing," *21st Annual Symposium Proceedings*, Society of Flight Test Engineers, Lancaster, CA, 1990, pp. 3.4-1–3.4-12.

[3]Cunningham, A. M., Jr., and Meijer, J. J., "Semi-Empirical Unsteady Aerodynamics for Modeling Aircraft Limit Cycle Oscillations and Other Non-Linear Aeroelastic Problems," *Proceedings of the International Forum on Aeroelasticity and Structural Dynamics*, Vol. 2, Royal Aeronautical Society, London, 1995, pp. 74.1–74.14.

[4]Meijer, J. J., and Cunningham, A. M., Jr., "A Semi-Empirical Unsteady Nonlinear Aerodynamic Model to Predict Transonic LCO Characteristics of Fighter Aircraft," AIAA Paper 95-1340, April 1995.

[5]Chen, P. C., Sarhaddi, D., and Liu, D. D., "Limit-Cycle Oscillation Studies of a Fighter with External Stores," AIAA Paper 98-1727, April 1998.

[6]Denegri, C. M., Jr., and Cutchins, M. A., "Evaluation of Classical Flutter Analyses for the Prediction of Limit Cycle Oscillations," AIAA Paper 97-1021, April 1997.

[7]Denegri, C. M., Jr., "Limit Cycle Oscillation Flight Test Results of a Fighter with External Stores," *Journal of Aircraft*, Vol. 37, No. 5, 2000, pp. 761–769.

[8]Woolston, D. S., Runyan, H. L., and Andrews, R. E., "An Investigation of Effects of Certain Types of Structural Nonlinearities on Wing and Control Surface Flutter," *Journal of the Aeronautical Sciences*, Vol. 24, No. 1, 1957, pp. 57–63.

[9]Tang, D. M., Henry, J. K., and Dowell, E. H., "Limit Cycle Oscillations of Delta Wing Models in Subsonic Flow," *AIAA Journal*, Vol. 37, No. 11, 1999, pp. 1355–1362.

[10]Tang, D. M., Dowell, E. H., and Virgin, L. N., "Limit Cycle Behavior of an Airfoil with a Control Surface," *Journal of Fluids and Structures*, Vol. 12, No. 7, 1998, pp. 839–858.

[11]Haykin, S., *Neural Networks, A Comprehensive Foundation*, Macmillan, New York, 1994, pp. 138–179.

[12]Rosenblatt, F., *Principles of Neurodynamics*, Spartan, New York, 1961, p. 85.

[13]Albano, E., and Rodden, W. P., "A Doublet-Lattice Method for Calculating Lift Distributions on Oscillating Surfaces in Subsonic Flows," *AIAA Journal*, Vol. 7, No. 2, 1969, pp. 279–285.

[14]Desmarais, R. N., and Bennett, R. M., "An Automated Procedure for Computing Flutter Eigenvalues," *Journal of Aircraft*, Vol. 11, No. 2, 1974, pp. 75–80.

[15]Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., "Leavenberg-Marquardt Method," *Numerical Recipes in FORTRAN*, 2nd ed., Cambridge Univ. Press, New York, 1992, pp. 678–683.