

Comparison of Static and Dynamic Neural Networks for Limit Cycle Oscillation Prediction

Michael R. Johnson* and Charles M. Denegri Jr.†

Air Force SEEK EAGLE Office, Eglin Air Force Base, Florida 32542-6865

A dynamic artificial neural network in the form of a multilayer perceptron with a delayed recurrent feedback connection is investigated to determine its ability to predict linear and nonlinear flutter response characteristics. Flight-test results show that limit cycle oscillation (LCO) response characteristics are strongly dependent on Mach number in the transonic flight regime. This effect is also evident in the classical transonic small-disturbance theory governing equations. A dynamic network is considered in order to examine the effects of sequential Mach-number dependence on the network's predictive capability. The architecture of a dynamic network allows for modeling data dependent on a sequentially or linearly increasing parameter (usually time, but in this case Mach number). The predictive capabilities are compared to those of a static artificial neural network. 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 dynamic network was successful in predicting the aeroelastic oscillation frequency and amplitude responses over a range of Mach numbers for two of the test cases. The dynamic network showed slightly better correlation to flight-test results for the typical LCO test case but slightly worse correlation for the flutter case. Predictions for the nontypical LCO test case were not good for either network.

Nomenclature

g	=	artificial structural damping
K	=	modal stiffness matrix
KCAS	=	knots calibrated airspeed
k	=	stiffness matrix
M	=	modal mass matrix
m	=	mass matrix
Q	=	generalized aerodynamic force matrix
q_{linear}	=	linear flutter analysis displacement
q_{measured}	=	flight-measured displacement
$q_{\text{nonlinear}}$	=	displacement from nonlinear sources
r	=	linear flutter analysis modal displacement
V	=	flutter solution velocity
$\Delta\omega(V)$	=	velocity-dependent linear frequency variation
$\Delta\omega_{\text{nonlinear}}$	=	frequency variation caused by nonlinear effects
ρ	=	air density
ϕ, ϕ_i	=	free-vibration mode shape
$\omega_i, \omega_{\text{natural}}$	=	natural frequency
$\omega_{\text{linear}}, \omega$	=	linear analysis flutter frequency
ω_{measured}	=	flight-measured frequency

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 papers 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 papers 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, whereas 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.

Received 2 February 2001; presented as Paper 2001-1289 at the AIAA/ASME/ASCE/AHS/ASC 42nd Structures, Structural Dynamics, and Materials Conference, Seattle, WA, 13–16 April 2001; revision received 25 July 2002; accepted for publication 18 August 2002. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0021-8669/03 \$10.00 in correspondence with the CCC.

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

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

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 due to 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 oscillation 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 (caused by 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 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). Previous studies⁸ showed the feasibility of using a static artificial neural network (static ANN) to map the results of flight tests into an intelligent system that infers the results of slightly differing configurations before any test missions are flown. The present work examines an extension to that approach by evaluating the capabilities of a dynamic artificial neural network (dynamic ANN) in the form of a multilayer perceptron with a delayed recurrent feedback connection. Flight-test results show that LCO response characteristics are strongly dependent on Mach number in the transonic flight regime. This effect is also evident in the classical transonic small-disturbance theory governing equations. The architecture of a dynamic network allows for modeling data dependent on a sequentially or linearly increasing parameter (usually time, but in this case equally effective for Mach number). Thus, the effects of sequential Mach-number dependence on the network's predictive capability can be examined.

In subsequent sections the use of the dynamic ANN for predicting LCO on a fighter aircraft with external stores is evaluated. The general analysis concept is that of constructing an artificial neural network based upon empirical inputs from linear flutter analysis and LCO flight tests. The network is then used 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, 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 simply 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.




Fig. 1a Single network node.




Fig. 1b Simple two-node network.




Fig. 1c Network nonlinear functions.

Neural-Network Fundamentals

The concept of the artificial neural network (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 interpolation⁹ within a problem space and can be used as a tool for the prediction of nonlinear states beyond the problem space bounds.

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.¹⁰ The term has come to be used interchangeably with "node," however. A node is illustrated in Fig. 1a. 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. 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.⁹ Output layers are more likely to be linear, either limited by upper or lower bounds or both, or unlimited. Figure 1b shows a small, two-input, single-output ANN.

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. Static networks are an obvious choice for modeling finite impulse response systems. If the system being modeled has a transfer function that has infinite impulse response (IIR) characteristics, a dynamic network might 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 to its desired response. Backpropagation is the process by which errors between the network response and the desired response are propagated back through the various layers of the network. A detailed description of backpropagation can be found in Ref. 9.

Table 1 External store configurations

Configuration	Station 1 wing-tip missile		Station 2 under-wing missile		Station 3 weapon		Station 4 fuel tank	Category
	Suspension equipment	Store	Suspension equipment	Store	Suspension equipment	Store	Fuel state	
2	Launcher A	None	Launcher A	None	Launcher C	Missile 3	No tank	Flutter
3	Launcher A	None	None	None	Launcher C	Missile 3	Empty	Flutter
7	Launcher A	None	Launcher A	None	Launcher C	Missile 3	Empty	Flutter
8 ^a	Launcher A	None	None	None	Launcher C	Missile 3	No tank	Flutter
1	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 3	$\frac{1}{4}$ full	Typical LCO
4	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 1	$\frac{1}{4}$ full	Typical LCO
6	Launcher A	None	Launcher A	Missile 1	Launcher C	None	$\frac{1}{2}$ full	Typical LCO
12	Launcher B	None	Launcher A	Missile 1	Launcher C	Missile 3	$\frac{1}{4}$ full	Typical LCO
13 ^a	Launcher A	None	Launcher A	Missile 2	Launcher C	Missile 3	Empty	Typical LCO
16	Launcher B	None	Launcher A	Missile 1	Launcher C	None	$\frac{1}{2}$ full	Typical LCO
18	Launcher B	None	Launcher A	Missile 2	Launcher C	Missile 3	Empty	Typical LCO
5	Launcher A	None	Launcher A	Missile 1	Launcher C	Missile 1	Empty	Nontypical LCO
15	Launcher A	None	Launcher A	Missile 2	Launcher C	Missile 1	Empty	Nontypical LCO
17 ^a	Launcher B	None	Launcher A	Missile 2	Launcher C	Missile 1	Empty	Nontypical LCO

^aNetwork test configuration.

Table 2 Flight-test response amplitude^a at 5000-ft altitude

Configuration	Mach number									Onset KCAS	Frequency, Hz	Category
	0.80	0.85	0.90	0.91	0.92	0.93	0.94	0.95	0.98			
2	0.0	0.0	0.0	N ^b	N	2.5	4.0	— ^c	—	569	9.40	Flutter
3	0.0	0.0	0.0	N	N	N	4.5	—	—	575	9.40	Flutter
7	0.0	0.0	0.0	N	N	N	N	0.5	—	582	9.20	Flutter
8 ^d	0.0	0.0	1.0	N	2.0	—	—	—	—	550	9.50	Flutter
1	0.0	0.5	1.5	N	N	N	N	2.5	—	519	6.60	Typical LCO
4	0.0	1.0	2.0	3.0	—	—	—	—	—	519	6.80	Typical LCO
6	0.0	0.0	0.5	N	N	N	N	1.0	1.5	550	6.90	Typical LCO
12	0.0	1.0	2.0	N	N	N	N	2.5	—	519	6.80	Typical LCO
13 ^d	0.0	1.0	3.0	4.0	—	—	—	—	—	519	7.80	Typical LCO
16	0.0	0.0	0.0	N	N	N	N	0.5	—	582	7.00	Typical LCO
18	0.0	1.0	2.5	3.0	—	—	—	—	—	519	8.10	Typical LCO
5	0.0	0.0	1.5	N	N	N	N	0.0	—	550	7.00	Nontypical LCO
15	0.0	0.5	1.0	N	N	2.0	N	0.0	—	519	8.10	Nontypical LCO
17 ^d	0.0	0.0	0.5	0.0	—	—	—	—	—	550	8.20	Nontypical LCO

^aAmplitude 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.

^bN-no response data explicitly measured.

^c— indicates no test data acquired.

^dNetwork test configuration.

LCO Characteristics

This section presents the flight test and linear analysis results of the external store configurations chosen for this study. Guidelines for understanding and interpreting these results are also discussed.

Reference 7 describes three categories of aeroelastic response behavior seen on fighter aircraft. These categories are described as 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 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 might be present only in a limited portion of the flight envelope.

Fourteen external store configurations are used for the present study and are presented in Table 1. This mix of configurations is typical for a certification flight-test program on fighter aircraft. These configurations have a variety of stores and suspension equipment including three different sets of missile launchers, missiles, and fuel quantities in the tanks. These configurations are diverse in their weapon carriage combinations, but they have similarities within each LCO response category. For example, all of the configurations that exhibited flutter response characteristics (configurations 2, 3,

7, and 8) are carrying Missile 3 on station 3 and have no missile present on station 2.

Flight-test results for this group of store configurations are presented in Table 2. The flight-test results are for level flight at 5000-ft 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 might 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.

As discussed in Ref. 7, the flight-test results for a particular configuration define its LCO response category. From the results presented in Table 2, it is seen that each configuration exhibits one of the three categories discussed earlier of aeroelastic response behavior (flutter, typical LCO, or nontypical LCO). The similarities in the store carriage configurations mentioned earlier lead to similarities in the instability oscillation frequencies for each category. The flutter configurations showed frequencies greater than 9 Hz, whereas the frequencies of the LCO configurations were all below 8.2 Hz.

Linear flutter analyses results for each flight-test configuration are presented in Table 3. These analyses are not matched analyses but merely worst-case "screening" analyses. In this manner all analyses

Table 3 Linear flutter analyses, 0.95 Mach, sea level

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

^a See Fig. 3 for node line representation of mode shapes.


^b First wing bending.

^c First wing torsion.

^d Network test configuration.

^e Forward wing bending.

^f Aft wing bending.

**Fig. 2** Flutter analysis model composed of doublet-lattice aerodynamics and lumped mass structure (all dimensions in inches).

are performed using sea-level density and 0.95 Mach doublet-lattice method¹¹ aerodynamic influence coefficients. (Note: The authors acknowledge that the use of 0.95 Mach for the doublet-lattice method stretches the applicability of the underlying theory. However, for this type of analysis reasonable correlation with the flight-measured LCO onset speeds has been established. These analyses are used here for consistency with conventional industry practice.) The free vibration analyses are performed for a half-airplane model using a matrix iteration method. The first 16 antisymmetric flexible modes are retained for the flutter solution. The flutter equations are solved using the Laguerre iteration method,¹² which is a variation of the classical k-method. The structural and aerodynamic models used for these analyses are shown in Fig. 2.

Similar to the flight-test results, the flutter analyses show similarities for the configurations within each LCO response category. The most consistent feature is that the flutter and typical LCO configurations each have the same combinations of modes in their respective unstable mechanisms. Also of interest is that the nontypical LCO configurations have these same modes present in their mechanisms, but the order of their frequencies and coupling are different from the other categories. Node line representations of the linear flutter mechanism modes are shown in Fig. 3.

The analytical flutter speed is considered to be the velocity at which the modal damping curve of the known aeroelastically sensitive mode for the particular configuration crosses from stable to unstable (0% damping). It is considered to be directly comparable to the lowest airspeed at which self-sustained oscillations, either LCO or flutter, are encountered in flight.

Comparing the flutter analyses (Table 3) to the flight-test results (Table 2), it is seen that the predicted speeds for the flutter configurations are much higher than the actual instability onset speed. The predicted speeds for the typical LCO configurations are more in line with the flight-test results, but there is no consistent correlation to the rate of oscillation amplitude increase. For the nontypical LCO cases the speeds for two of the configurations are reasonable, but there is no indication that the oscillations will cease at higher speeds for any of the configurations.

Fundamental Assumptions

The discussion in the LCO characteristics section showed a fundamental connection between linear flutter analyses and LCO flight-test results. This discussion also highlighted the weaknesses and limitations of using linear flutter analyses to predict LCO characteristics. From this, it is evident that linear flutter analyses can be




Fig. 3 Linear flutter analysis mode shapes (see Table 3 for frequencies). (Note: Forward and aft wing bending modes have both pitch and plunge components at the wing tip. First wing bending has only plunge. First wing torsion has only pitch.)

used for the ANN analysis, but to do so some assumptions regarding LCO are required. 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 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 being 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 might only serve to limit the amplitude of the resulting oscillations and do not change the flutter speed of the aircraft configuration.¹³ This assumption is supported by observations during flight tests of LCO-prone fighter aircraft, which show that the flutter speed is relatively independent of the strength of the disturbances used to excite the aeroelastic system.

The third fundamental assumption is that the physical nonlinearities that cause LCO in a new configuration must be adequately represented in the historical data and these nonlinearities are inherent in the wing structure or aerodynamics. The nonlinearities of the new configuration are considered to be essentially the same as for the flight-test configurations used in training the artificial neural network, and these nonlinearities are assumed to not vary from one configuration to another. That is, it is assumed that the store configuration can change the flutter speed or LCO onset speed, but not the nonlinearity.

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.

Analysis Parameters

With the fundamental assumptions, LCO characteristics, and neural network fundamentals in mind, we can now examine the underlying aeroelastic theory for determination of appropriate input and

output parameters for the neural network. For ease of relating the analysis to the flight-test predictions, it is desired to use physically interpretable parameters for the network input, output, and training. The flight-measured displacements meet this criterion and are a practical choice for a network output.

Assuming that the measured displacements are composed of a linear and nonlinear part, then

$$q_{\text{measured}} = q_{\text{linear}} + q_{\text{nonlinear}}$$

The measured displacements are known from the flight-test results, and the linear displacements can be obtained from classical linear flutter theory making use of the modal transformation

$$\{q_{\text{linear}}\} = [\phi]\{r\}$$

The mode shape matrix necessary to complete this transformation is obtained from solution of the free-vibration equation

$$(-\omega_i^2[m] + [k])\{\phi_i\} = 0$$

and the modal displacements are obtained from solution of the linear flutter equation.

$$[-\omega^2 M + (1 + ig)K - \frac{1}{2}\rho V^2 Q]\{r\} = 0$$

In a similar fashion the flight-measured frequencies can be decomposed into a linear component and a variation caused by the nonlinear effects

$$\omega_{\text{measured}} = \omega_{\text{linear}} + \Delta\omega_{\text{nonlinear}}$$

where

$$\omega_{\text{linear}} = \omega_{\text{natural}} + \Delta\omega(V)$$

All of the parameters necessary to complete the linear analysis are available for training the neural network either directly or indirectly. Table 4 lists the important parameters, their sources, and how they are directly or indirectly represented in the neural network solution process. This still leaves the nonlinear displacement component as an unknown, which is essentially obtained from the neural network during the training process.

With useful analysis parameters now defined, 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 were presented in Table 2. These configurations exhibit characteristics that are representative of the broad spectrum of flutter and

Table 4 Theoretical parameters for ANN analysis

Parameter	Source	Type
q_{measured}	Flight-measured response	Direct
ω_{measured}	Flight-measured frequency	Direct
q_{linear}	Linear flutter analysis	Indirect
$q_{\text{nonlinear}}$	Neural network	Indirect
ϕ, ϕ_i	Free-vibration analysis	Direct
r	Linear flutter analysis	Indirect
$\omega_{\text{linear}}, \omega$	Linear flutter analysis	Direct
$\omega_i, \omega_{\text{natural}}$	Free-vibration analysis	Direct
$\Delta\omega(V)$	Linear flutter analysis	Indirect
M	Linear flutter analysis, store carriage configuration, free-vibration modes and frequencies	Indirect
K	Linear flutter analysis, store carriage configuration, free-vibration modes and frequencies	Indirect
Q	Linear flutter analysis, Mach no., store carriage configuration, free-vibration modes and frequencies	Indirect
g	Linear flutter analysis	Direct
V	Linear flutter analysis	Direct
ρ	Flight altitude	Indirect

LCO responses encountered by fighter aircraft with external stores. Reference 6 categorizes these characteristics and showed a distinct correlation between the modes comprising the predicted linear flutter mechanism and the flutter or LCO behavior. Essentially, it was shown that different linear analysis flutter mechanisms correlated to different aeroelastic responses in flight. Based on these results, the structural 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 (Table 3 and Fig. 3). The aerodynamic characteristics are represented to the ANN by the linear analysis flutter speed and frequency and by a quantized representation of the wingtip 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. A complete list of the input and output parameters are shown in Table 5.

Some network inputs are quantized values of descriptive information. Because the network requires numerical quantities as input,

all of the descriptive data are represented as integers. Table 6 shows the quantization values used for the input descriptive data.

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. In this manner the oscillation amplitude and frequency trends with respect to Mach number can be evaluated and compared to the flight-test results.

Network Design

A static neural network in the form of a multilayer perceptron (MLP) was chosen for a previous study⁸ because of its ease of design and evaluation, as well as its simplicity. The static network was designed to simply feed forward input sets one at a time and provide a prediction. By its nature there was no inherent knowledge of past inputs or outputs. Therefore, the problem was simply one of functional representation of the training data. For the present work a dynamic neural network in the form of a MLP with a delayed recurrent feedback connection is designed in an attempt to further refine the predictive capability of the network. It is felt that by applying a feedback loop from the output back to the input the network would be better able to predict the next output because it would now have knowledge of the preceding prediction. Specifically, the network is comprised of an input layer feeding into a hidden layer, followed by another, smaller hidden layer, in a manner similar to earlier static networks. The output is a final, linear layer that provides frequency and amplitude of the predicted oscillation. This output is also fed back to the input layer through its own weight structure as further input to the network.

Transonic small-disturbance theory¹⁴ indicates a dependence on Mach number and practical observations of the LCO flight characteristics support this. The feedback connection in the dynamic ANN is an attempt to model this sequential Mach-number dependence. In other words, consider the test data not as an explicit time sequence but rather a set with a dependence on (or knowledge of) prior events (such as a linearly increasing function of Mach number). This dependence is the motivation for considering dynamic ANNs in the present work. The basic premise being that if LCO or flutter occurred at a lower Mach number it is likely to occur at a higher Mach number, also. From a physical perspective, if LCO is strongly Mach dependent, then allowing the network to connect predicted behavior at prior Mach numbers will refine its prediction for the current Mach number. This attempts to give the network a broader view of the aeroelastic behavior potential of a particular configuration. In effect, the network now has memory of past events.

Details of the dynamic network are shown in Fig. 4. The network 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. The outputs of the hyperbolic tangent nodes are then weighted again, summed, and fed into another hyperbolic tangent layer. The process is repeated again, this time through saturated linear output nodes. The output is then fed back to the first hidden layer through a delay, and multiplied by a feedback weight matrix, forming a recursive loop in the network. Both the hyperbolic tangent and saturated linear functions have the property of limiting outputs to the range $[-1, 1]$. These functions are shown in Fig. 1c.

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 Levenberg–Marquardt algorithm used in the previous work⁸ was not applicable for training recurrent networks. Methods based on Newton's method of gradient descent have been shown to perform more consistently in dynamic networks. For this study a quasi-Newton method using an approximate Hessian matrix was used to train the dynamic networks, where the weight update is computed as a function of the gradient. This method did not converge to a solution as fast as Levenberg–Marquardt, but gave more consistent results during multiple training sessions.

Typically, the static ANN would converge in seven to 25 epochs using Levenberg–Marquardt, whereas the dynamic ANN required

Table 5 Network input and output parameters

Parameter	Data type	I/O
Mach number	Real number	Input
Sta. 1 launcher	Integer representation of name	Input
Sta. 2 launcher	Integer representation of name	Input
Sta. 2 store	Integer representation of name	Input
Sta. 3 store	Integer representation of name	Input
Sta. 4 fuel state	Integer representation of name	Input
Flutter speed (0% damping)	Real number	Input
Flutter speed (1% damping)	Real number	Input
Flutter frequency	Real number	Input
Unstable flutter mechanism mode shape	Integer representation of name	Input
Unstable mode natural frequency	Real number	Input
Coupled flutter mechanism mode shape	Integer representation of name	Input
Coupled mode natural frequency	Real number	Input
Oscillation amplitude	Real number	Output
Oscillation frequency	Real number	Output

Table 6 Quantization of input descriptors

Input parameter	Descriptor	Value
Sta. 1 launcher	Launcher A	1
	Launcher B	2
Sta. 2 launcher	Launcher A	1
	No launcher	0
Sta. 2 store	Missile 1	1
	Missile 2	2
	No missile	0
Sta. 3 store	Missile 1	1
	Missile 3	3
	No missile	0
Sta. 4 fuel state	$\frac{1}{2}$ full tank	3
	$\frac{1}{4}$ full tank	2
	Empty tank	1
	No tank	0
Unstable mode	Forward wing bending	0
	First wing torsion	1
	Aft wing bending	2
	First wing bending	3
Coupled mode	Forward wing bending	0
	First wing torsion	1
	Aft wing bending	2
	First wing bending	3

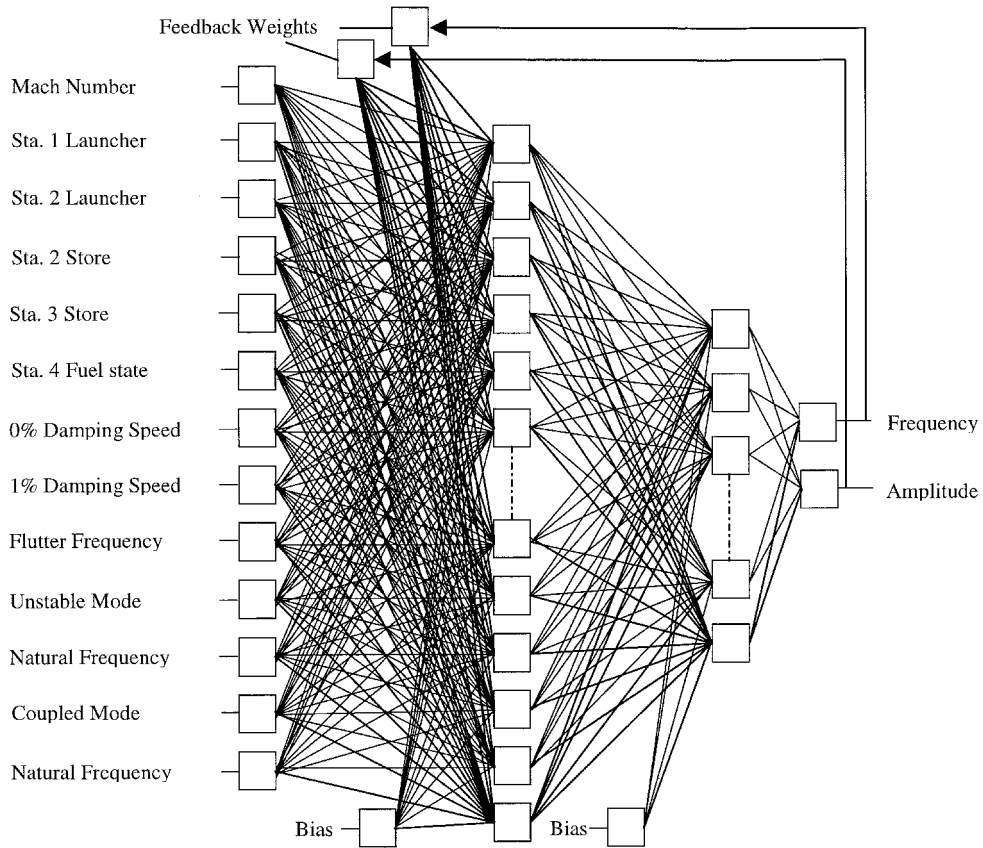


Fig. 4 Network diagram.




Fig. 5 Flight-test response characteristics, flutter subcases. (Note: Flutter subcases shown separately here for clarity. Network is trained using all subcases except configurations 8, 13, and 17.)




Fig. 6 Flight-test response characteristics, typical LCO subcases. (Note: Typical LCO subcases shown separately here for clarity. Network is trained using all subcases except configurations 8, 13, and 17.)

700–2000 epochs. The tradeoff, however, is that the quasi-Newton method requires many fewer calculations per epoch, so that the total training time was not significantly longer. An epoch is defined as one presentation of the training data to the network during the network training process. All networks considered used the data shown in Tables 1–3 and Figs. 5–7. Thirteen inputs were given for each output frequency and amplitude vs Mach combination (Fig. 4 and Table 5). 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 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. 8–10) 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. 5–7. 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 define the aeroelastic response category of the test configuration results.

Several networks were designed, trained, and evaluated in order to find the optimal configuration. Single- and double-layer networks were evaluated. As with the static network in Ref. 8, a double-layer




Fig. 7 Flight-test characteristics, nontypical LCO subcases. (Note: Nontypical LCO subcases shown separately here for clarity. Network is trained using all subcases except configurations 8, 13, and 17.)




Fig. 9 Typical LCO test case. (Note: Neural networks evaluated at 0.01 Mach increments in order to show analysis trends.)

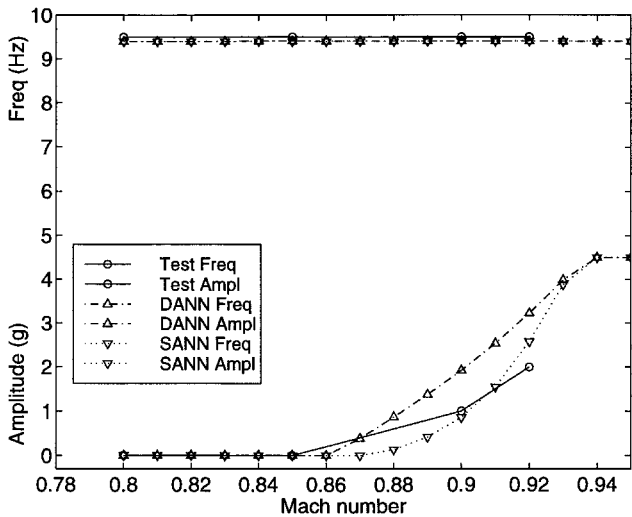


Fig. 8 Flutter test case. (Note: Neural networks evaluated at 0.01 Mach increments in order to show analysis trends.)




Fig. 10 Nontypical LCO test case. (Note: Neural networks evaluated at 0.01 Mach increments in order to show analysis trends.)

network was chosen because it allowed more accurate modeling of the nonlinear characteristics of the solution space.

Consideration was given to network size, but only to the extent that accurate functional representation could be achieved. Because the size of the required static ANN had been previously determined,⁸ the search for an acceptable network configuration began with networks of comparable size or larger. Minimization of the size of the network was not the focus of this work and was not considered.

During the course of training networks the quality of the solution to the training data was observed by plotting the network approximations to the training data set after training was accomplished. It was an easy task, therefore, to determine if the network was a reasonable, but not overfit functional approximation to the system.

The final design scheme selected for the dynamic MLP contains 39 hyperbolic tangent nodes in the first hidden layer and 12 in the second, with two saturated linear nodes as output (39×12). This configuration gave good generalization over the solution space, while learning the training data well. The outputs were the combined oscillation frequency in hertz, and oscillation amplitude measured in units of gravitational acceleration (g) at the forward end of the wingtip launcher. All input data were normalized to the range $[-1, 1]$. The output was also contained in the range $[-1, 1]$ because of the saturated linear output nodes. All output data were 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, the linear analysis 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. 5–7) is shown in Figs. 8–10 as two two-dimensional subplots of frequency and amplitude vs. Mach number. From a two-dimensional perspective the test data in Figs. 8–10 appear to be simple polynomials (or straight lines). However, keep in mind that 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. When evaluating network predictions, it is also important to note that the networks are trained using constant amplitude response data. Therefore, 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. So, it is possible for an oscillatory frequency to be associated with zero-amplitude response. The frequencies associated with zero-amplitude responses are the dominant frequencies of the damped responses.

Several Mach numbers were examined for each set of input conditions, and responses were noted. The classification of the predicted response is by direct comparison to the flight-test data and 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 increased up to a particular Mach number and then decreased for higher Mach numbers.

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 both the static and dynamic networks are compared in Figs. 8–10. The measured test case data are available only at 0.80, 0.85, 0.90, 0.91, and 0.92 Mach. However, the networks are evaluated at 0.01 Mach resolution from 0.80 to 0.95 Mach in order to more clearly observe the network functional characteristics.

Comparisons of static and dynamic networks for the flutter test case are shown in Fig. 8. The static network (SANN) shows very good correlation to the test data at all measured test points. This network shows a subtle increase in oscillation amplitude beginning at 0.88 Mach. The amplitude curve slope exponentially increases at higher Mach numbers up to 0.94 Mach. The amplitude response predicted by the static network is consistent with that seen in the flight-measured data. By comparison, the dynamic network (DANN) shows a sudden upturn in response amplitude at 0.87 Mach. The amplitude curve slope increases almost linearly up to 0.94 Mach. The slope of the dynamic network response is not as steep as that of the static network. Although the onset Mach numbers and peak response amplitudes are very similar for the two networks, in practice, the dynamic network results appear to be more indicative of LCO behavior rather than flutter. In other words, the static network indicates flutter while the dynamic network indicates LCO. The dynamic network shows a progressive increase in amplitude with respect to Mach number rather than a sudden increase, which would be indicative of flutter. The computed response frequencies for both networks showed very good correlation to the measured data.

For the typical LCO test case (Fig. 9) the SANN amplitude tracks the measured response very well at the measurement test points. However, the functional behavior of the static network shows a greater disparity with the general trend of the flight-test data. The relatively steep slope in the 0.83–0.88 Mach range indicates a rapid increase to higher level LCO that is not seen in the flight-test data. Also, at the last measured test point at 0.91 Mach the static network output shows a trend of leveling off while the measured response amplitude continues to increase. By comparison, the DANN also tracks the measured response very well at the measurement test points but shows much better agreement with the typical LCO response trend of the measured data. As was observed in the flutter case, it is seen that both networks predicted the frequencies of oscillation nearly exactly.

Table 7 Distribution of network training and testing points

Response category	Training pts.	Testing pts.	Train-test ratio	Training configurations
Flutter	13	4	3.25	3
Typical LCO	25	4	6.25	6
Nontypical LCO	9	4	2.25	2
Total	47	12	3.92	11

As with the previous work,⁸ the nontypical LCO case was the most revealing of the three test cases in that it again challenged the ANN beyond its capabilities with the given data set. As shown in Fig. 10, the static ANN amplitude tracked the flight response well up to 0.91 Mach. There, the static network indicated a slight increase in amplitude while the flight data showed a decrease. The dynamic ANN had trouble tracking the amplitude from 0.83 Mach and beyond where its response was more indicative of typical rather than nontypical LCO. Neither network tracked frequency as well as for the preceding cases.

Conclusions

The feasibility study⁸ concluded that the static ANN was successful considering the small data set used for training and the limitations of the static network itself. The present work shows that the flutter test case is adequately predicted, and the typical LCO test case is more accurately predicted with the dynamic network. A limitation of the dynamic network is suggested in the nontypical LCO case, where the network had difficulty predicting the decreasing change in oscillation amplitude that would be indicative of this category of response behavior.

For the flutter case the dynamic ANN showed a relatively steep increase to high response amplitudes as Mach number was increased. However, the slope of the response curve was not as steep as that of the static network and could be misinterpreted in practice as indicating typical LCO rather than flutter. For the typical LCO case the general behavior of the dynamic ANN tracked that of the flight-test data showing a progressive increase in response amplitude as Mach number was increased. For very limited data sets, such as the nontypical LCO case, the dynamic network performs no better than the static ANN does. The dynamic ANN had trouble tracking the amplitude for the nontypical LCO case and showed response that was more indicative of typical LCO.

Examination of these results along with a closer look at both network types provides some insight into modeling approaches. The static ANN, whose strength is its ability to map an input pattern to an output pattern, many times gave better, or at least as good results as the dynamic ANN, whose primary use is tracking sequential data. The underlying physics, however, indicates dependence on prior response characteristics (in this case, Mach number), implying that use of a dynamic network would be more appropriate. But the dynamic network only produced better predictions for the typical LCO case, which by definition shows progressive increases in oscillation amplitude with increasing Mach number. This infers, then, that flutter response characteristics are sufficiently different from typical LCO as to require a different modeling paradigm. This inference cannot yet be extended to the nontypical LCO response because of the small number of samples available for training.

Examining the functional form of the flight-test phenomenon being predicted (Figs. 5–7) and the quantity of data available for training and testing the networks (Table 7), it is observed that the data set for the nontypical LCO case is much smaller than for the other two response classes. One of the fundamental assumptions was that sufficient data must exist from which the network can infer a reasonable prediction of a trend in the physical phenomenon. The fact that the nontypical LCO results do not show good correlation to the flight-test results indicates that this assumption is critical to obtaining reasonable and viable predictions from the neural network approach and offers some quantitative bounds on what constitutes “sufficient data.”

It is concluded that the optimal ANN modeling of flutter behavior is a series of patterns in Mach and response amplitude space, best approached using a static ANN. And the optimal ANN modeling of typical LCO, and most likely nontypical LCO, is a recursive system, best approached using a dynamic ANN dependent not only on a sequential set of inputs but the preceding results as well. This suggests that in order to utilize a single network architecture for modeling all three types of response characteristics the optimal network type must be dynamic. Note that this does not contradict our conclusion for flutter response characteristics because of the flexibility of a dynamic network, which is structured in such a way as to optimize its use of the feedback connections, thereby drawing on them only when necessary, thus emulating a static network when required.

Given a sufficient historical database of flight-test results, the ANN approach for predicting LCO characteristics can provide additional information that is useful in the certification process beyond what is available from linear flutter analyses. Linear analyses adequately identify store configurations that are flutter and LCO sensitive, and these analyses give a good indication of the instability frequency. But, linear flutter analyses do not accurately identify the instability onset speed or the amplitude of the unstable response. By comparison, the ANN approach provides predictions of amplitude, Mach, and altitude trends similar to those obtained from flight testing. With improvement in the reliability of the predictions, the availability of this additional data will allow certifications that rely less heavily on flight testing.

References

- ¹Bunton, 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.
- ²Norton, 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.
- ³Cunningham, 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.
- ⁴Meijer, 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.
- ⁵Chen, P. C., Sarhaddi, D., and Liu, D. D., "Limit-Cycle Oscillation Studies of a Fighter with External Stores," AIAA Paper 98-1727, April 1998.
- ⁶Denegri, 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.
- ⁷Denegri, 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.
- ⁸Denegri, C. M., Jr., and Johnson, M. R., "Limit Cycle Oscillation Prediction Using Artificial Neural Networks," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 5, 2001, pp. 887–895.
- ⁹Haykin, S., *Neural Networks, A Comprehensive Foundation*, Macmillan, New York, 1994, pp. 138–179.
- ¹⁰Rosenblatt, F., *Principles of Neurodynamics*, Spartan Press, Washington, DC, 1961, p. 85.
- ¹¹Albano, 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.
- ¹²Desmarais, Robert N., and Bennett, Robert M., "An Automated Procedure for Computing Flutter Eigenvalues," *Journal of Aircraft*, Vol. 11, No. 2, 1974, pp. 75–80.
- ¹³Woolston, 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.
- ¹⁴Dowell, E. H., Curtiss, H. C., Jr., Scanlan, R. H., and Sisto, F., *A Modern Course in Aeroelasticity*, Kluwer Academic, Norwell, MA, 1989, pp. 239–241.