Modeling and Detection of Limit-Cycle Oscillations Using Adaptable Linear Models

Michael R. Johnson[*]
U.S. Air Force SEEK EAGLE Office, Eglin Air Force Base, Florida
and
Jose C. Principe[†]
University of Florida, Gainesville, Florida 32611

A method for modeling the flutter response of a thin winged aircraft is presented. A hybrid physical-adaptive modeling framework is proposed to separate the autoregressive and moving average flutter components. Adaptive oscillators set at the structural free-vibration modal frequencies of the wing represent the structure (the autoregressive component). The moving average filters represent signal changes caused by aerodynamic forces encountered during flight. Connected in series, these modules form a hybrid model for wing flutter under changing flight conditions. The moving average component is trained to predict the signal produced by an accelerometer at the wing tip. The residual is segmented using an analysis of variance. The resulting family of linear moving average models provides a synthesis of the wing's response over time. Our analysis shows that this modeling paradigm performs well with data taken at increasing Mach numbers in level flight. Network parameters are shown to correlate linearly with Mach. Network components can themselves be predicted as functions of Mach number, allowing the accelerometer signals to be predicted with a high degree of accuracy. The paradigm is further validated by using an adapted model with a separate set of data from a different but similar flight condition.

Introduction

ONE of the most interesting problems associated with flight is the interaction of aerodynamic, inertial, and structural forces that result in flutter.[1] Recent engineering has noted and studied a subclass of flutter dealing with the phenomenon of limit-cycle oscillations (LCO).[2] These oscillations are a persistent problem on modern fighter aircraft with thin, flexible wings, and new ways to describe, model, and simulate these responses are always being sought.

Modern flight-test organizations now have the advantage of large amounts of data obtained from onboard instrumentation systems. This increase in data availability expands the possible flutter and LCO analysis options. We address the study of LCO, not from a traditional aeroelastic perspective, but as a challenge to be answered using adaptive signal-processing techniques. The data contained in the instrumentation signals can provide fresh insight into the structural behavior of the aircraft. Using a dynamic artificial-neural-network architecture, this work goes beyond evaluating the LCO response by simple observation of envelop characteristics to a process that gives a clearer understanding of the modal construct of the signal and the response of the structure itself.

The particular variety of LCO we will focus on is that observed in the wing's behavior during transonic flight. It is generally accepted that LCO arises from the interaction of structural and aerodynamic forces acting on the affected aircraft component, and that this interaction is nonlinear in nature.[3] The contribution of forces, as well as their interaction, has been the topic of much research and debate. Classical flutter is a phenomenon that once initiated catastrophically diverges. Limit-cycle oscillation is a phenomenon described by Denegri and Dubben[4] as a variety of classical flutter, but differs in its tendency toward limited amplitude rather than diverging

oscillations. As such, we will treat flutter and LCO as the same phenomenon and assume the same background, theory, fundamental assumptions, model structure, and paradigm can be used for both.

Flutter analysis requires the determination of both the flutter speed and frequency. The analysis is, therefore, a double characteristic value problem, contrasted with vibration analysis that has only frequency as its single characteristic value. The methods of analysis are dictated by our limited knowledge of unsteady aerodynamic forces. As we gain knowledge, more exact solutions can be found. Presently, more extensive theoretical solutions are available for the case of steady-state harmonic motion because of the relative simplicity of the mathematical formulation, so that it is convenient to assume harmonic motion and use its methods for finding flutter speed and frequency. This is the premise of classical linear flutter analysis.[5]

But there is a large analytical gap between the current empirical modeling tools and their ability to match actual in-flight responses. Most techniques are linear and time invariant, whereas the actual in-flight response is highly nonlinear and time varying. If linear techniques prove insufficient, the traditional alternative is to use techniques that solve the equations of motion recursively such as computational fluid dynamics[6] (CFD).

The use of CFD has become more practical in recent years as computer speed and efficiency increases. But CFD is still very computationally intensive, and simplifications must be made.[7] Simulation of one aircraft configuration at one flight condition can take several hours. Because a single flight test for a given aircraft configuration can contain nearly 100 test conditions, it is easy to see that CFD is still far from being a practical certification tool.

Linear Flutter Analysis

The accepted linear flutter analysis governing equation for P dynamic mass model (DMM) points on a structure (a wing in our case) and n_m modes is[3]

$$m \frac{\mathrm{d}^2 q_v}{\mathrm{d}t^2} + C \frac{\mathrm{d}q_v}{\mathrm{d}t} + K q_v = A \qquad (1)$$

where m is the mass, a $P \times P$ diagonal matrix containing the mass for each DMM point; C is a $P \times P$ matrix of damping coefficients; K is stiffness, a $P \times P$ diagonal matrix containing the stiffness coefficient for each DMM point; q_v is wing vertical modal deflection, a

*Lead Engineer, 46 SK/SKA, 205 W. D. Avenue, Suite 348. Senior Member AIAA.
†Distinguished Professor of Electrical and Computer Engineering, Director, Computational NeuroEngineering Laboratory, EB 451, Building #33.

$P \times 1$ vector containing the superposition of the modal deflections at each DMM point; and A is an aerodynamic forcing function, a $P \times 1$ vector representing the total aerodynamic force at each DMM point.

If we make the assumption of simple, undamped harmonic motion in a direction perpendicular to the wing surface, the description of each DMM point can be reduced to independent spring-mass

for a 16-point wing analysis. Realizing that

$$\frac{A_i}{s^2 m_i + K_i} \xleftrightarrow{L^{-1}} \frac{A_i}{\sqrt{m_i K_i}} \sin\left(\sqrt{\frac{K_i}{m_i}}t\right)$$

the 16×16 system response for Eq. (1) is

$$h(t) = \begin{bmatrix} \frac{AIC_{1,1}}{\sqrt{m_1 K_1}} \sin\left(\sqrt{\frac{K_1}{m_1}}t\right) & \frac{AIC_{1,2}}{\sqrt{m_2 K_2}} \sin\left(\sqrt{\frac{K_2}{m_2}}t\right) & \cdots & \frac{AIC_{1,16}}{\sqrt{m_{16} K_{16}}} \sin\left(\sqrt{\frac{K_{16}}{m_{16}}}t\right) \\ \frac{AIC_{2,1}}{\sqrt{m_1 K_1}} \sin\left(\sqrt{\frac{K_1}{m_1}}t\right) & \frac{AIC_{2,2}}{\sqrt{m_2 K_2}} \sin\left(\sqrt{\frac{K_2}{m_2}}t\right) & \cdots & \frac{AIC_{2,16}}{\sqrt{m_{16} K_{16}}} \sin\left(\sqrt{\frac{K_{16}}{m_{16}}}t\right) \\ \vdots & \vdots & \ddots & \vdots \\ \frac{AIC_{16,1}}{\sqrt{m_1 K_1}} \sin\left(\sqrt{\frac{K_1}{m_1}}t\right) & \frac{AIC_{16,2}}{\sqrt{m_2 K_2}} \sin\left(\sqrt{\frac{K_2}{m_2}}t\right) & \cdots & \frac{AIC_{16,16}}{\sqrt{m_{16} K_{16}}} \sin\left(\sqrt{\frac{K_{16}}{m_{16}}}t\right) \end{bmatrix} \tag{7}$$

systems. The result is a simplification of mass and stiffness matrices so that the matrices become diagonal.

A structure such as the wing can be described as a multiple-input multiple-output (MIMO) system, where the resulting point deflections are outputs produced by a uniform initial displacement at each point. We know from examination of Eq. (1) that the system will have an infinite impulse response (IIR). The transfer function of the wing system under aerodynamic pressure given in Eq. (1) can be determined by its Laplace transform. Realizing that the vector q_v is a sum of modal displacements (we use a 16 DMM points and consider all 16 modes in this example to ease illustration)

$$q_v = \begin{bmatrix} q_{1,1} + q_{1,2} + \cdots + q_{1,16} \\ q_{2,1} + q_{2,2} + \cdots + q_{2,16} \\ \vdots \\ q_{16,1} + q_{16,2} + \cdots + q_{16,16} \end{bmatrix} \tag{2}$$

Next we describe q as a matrix of these individual modal displacements:

$$q = \begin{bmatrix} q_{1,1} & q_{1,2} & \cdots & q_{1,16} \\ q_{2,1} & q_{2,2} & \cdots & q_{2,16} \\ \vdots & \vdots & \ddots & \vdots \\ q_{16,1} & q_{16,2} & \cdots & q_{16,16} \end{bmatrix} \tag{3}$$

Beginning with Eq. (1) with $C = 0$ and m and K diagonal matrices, and replacing A with the product of a matrix of aerodynamic interference coefficients and q, $[AIC]q$, Eq. (1) and its Laplace transform become

$$m\frac{d^2 q}{dt^2} + Kq = [AIC]q \xleftrightarrow{L} s^2 m Q(s) + K Q(s) = [AIC]Q(s) \tag{4}$$

The MIMO transfer function for the wing under flutter can be written as

$$Q(s) = H(s) = [AIC][s^2 m + K]^{-1} \tag{5}$$

With our assumptions the matrix $[s^2 m + K]^{-1}$ is diagonal, and as such the inverse is easily calculated. The product $[AIC][s^2 m + K]^{-1}$ can be shown to be

$$H(s) = \begin{bmatrix} \frac{AIC_{1,1}}{s^2 m_1 + K_1} & \frac{AIC_{1,2}}{s^2 m_2 + K_2} & \cdots & \frac{AIC_{1,16}}{s^2 m_{16} + K_{16}} \\ \frac{AIC_{2,1}}{s^2 m_1 + K_1} & \frac{AIC_{2,2}}{s^2 m_2 + K_2} & \cdots & \frac{AIC_{2,16}}{s^2 m_{16} + K_{16}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{AIC_{16,1}}{s^2 m_1 + K_1} & \frac{AIC_{16,2}}{s^2 m_2 + K_2} & \cdots & \frac{AIC_{16,16}}{s^2 m_{16} + K_{16}} \end{bmatrix} \tag{6}$$

This matrix formulation reveals that the motion at each wing point P is a sum of its corresponding row, yielding

$$h_p(t) = \sum_{i=1}^{16} \frac{AIC_{p,i}}{\sqrt{m_i K_i}} \sin\left(\sqrt{\frac{K_i}{m_i}}t\right) \tag{8}$$

Notice that this impulse response represents a multiple input single-output system (MISO) and reflects the decoupling produced by the diagonal assumption of m and K. During flight, AIC is, in general, a time-varying matrix. However, Fung's premise of simple harmonic motion[5] suggests that flutter can be considered a time-invariant process for a short period of time. In this paper we will attempt to identify one of these MISO systems by using a realistic set of wing modes obtained from linear flutter analysis and utilizing one of the sensors placed in the instrumented wing of a flying aircraft to optimally determine the aerodynamic forces.

Modeling Approach

Equation (8) shows that the deflection of a DMM point on the wing is a linear combination of flutter modes. As such, a physical model that generates, in parallel, the modal signals can implement the modes, properly placing the poles of the system, then adjust each mode's amplitude as necessary. Signal generators, initialized to the structural free-vibration frequencies for optimal structural representation, can provide sinusoidal signals. Adaptation of each modal frequency will account for small variations in frequency. Adaptable gain filters will achieve modal amplitude scaling translating the aerodynamic forces. The model output is the sum of these filtered sinusoids. Further, by increasing the filter order, phase as well as gain for each mode can be adjusted. The initial MIMO autoregressive-moving-average (ARMA) model of Eq. (5) is then substituted by several MISO ARMA models of Eq. (8) at several locations on the wing, each of which can further be simplified as a cascade of two modules; one to act as signal generators and the other a bank of finite-impulse-response (FIR) filters. A conceptual diagram of the system is shown in Fig. 1.

A time-domain representation of the output of this MISO system is

$$y(t) = \sum_{m=1}^{16} a_m \sin(\omega_m t + \phi_m) \tag{9}$$
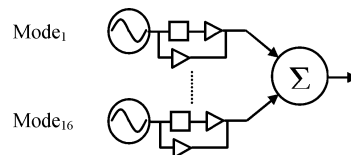

Fig. 1 Conceptual diagram of oscillator-driven FIR filter system.

where m is the mode number taken from the first 16 solutions of the linear flutter analysis, ω_m and ϕ_m is the frequencies and phases for each mode, and a_m is a set of amplitude parameters.

In this paper we will pursue a totally adaptive solution to set the parameters of Eq. (9) utilizing data collected from actual flight. The accelerometer sensors on the wings are considered the desired responses to optimally set the parameters using a system identification framework. What is interesting is that the use of the physical properties of the wing (modal analysis) and the simplifying assumptions of mode decoupling transform the difficult problem of adapting an ARMA model into the much simpler problem of adapting a set of FIR filters. Flutter is time varying, although exploitation of Fung's quasi-steady simplification allows us to consider it as locally stationary for short periods of time. Therefore we also have to include a tracking and segmenting scheme in the adaptive flutter model.

Adaptive Flutter Model Implementation

Oscillatory Module

A two-processing-element (PE) recurrent network, shown in Fig. 2, provides unit amplitude sinusoidal oscillations. This network is unique in that it has only one adaptable weight[8] ω_m used to directly adjust the frequency. The PEs to implement the oscillator are quasilinear (linear range but saturating at ± 1) to avoid instability during adaptation. The transfer function of each of the oscillators after adaptation is

$$H_o = \frac{\left\{ W_{1m} - [2\cos(\omega_m T)W_{1m} + W_{2m}]z^{-1} \right\}}{\left[1 - 2\cos(\omega_m T)z^{-1} + z^{-2} \right]} \quad (10)$$

As described in Ref. 8, all inputs are simply impulses at time $t = 0$, scaled by the input weights W_{1m} and W_{2m}. The system is simplified further with $W_{1m} = 0$ and $W_{2m} = -\sin(\omega_m T)$, yielding a sine wave with zero phase at the frequency ω_m. These adaptable oscillators act as the (autoregressive) AR portion of the MISO model. The poles can be adjusted to any desired value on the unit circle by propagating the error signals through the dual of the FIR filter as done in the backpropagation algorithm. Note however that because the oscillators are recurrent systems, backpropagation through time[9] (BPTT) is required to adapt the single free parameter.

Moving-Average (MA) Module

The design of the MA portion of the model, provided by FIR filters, is very straightforward. For this work a typical adaptable FIR filter will be implemented for each mode as

$$y(n) = W_0 x(n) + W_1 x(n-1) + \cdots + W_N x(n-N) \quad (11)$$

The weights $W_0, W_1, \ldots, W_N$, where N is the filter order, are adaptable. Because the purpose of the FIR filter is to adjust phase and gain of a sine wave, a first-order filter ($N = 1$) is sufficient. The final 16-mode architecture is shown in Fig. 3.

Now that the separate AR and MA models have been defined the transfer function for the model becomes

$$H_T(z) = \sum_{m=1}^{M} \frac{W_{i2m} z^{-1} \left(W_{0m} + W_{1m} z^{-1} \right)}{\left[1 - 2\cos(\omega_m T)z^{-1} + z^{-2} \right]} \quad (12)$$

Analysis of Variance for Signal-Change Detection

The only difficulty left is tracking the nonstationary nature of the aerodynamic forcing function. There are two basic approaches to
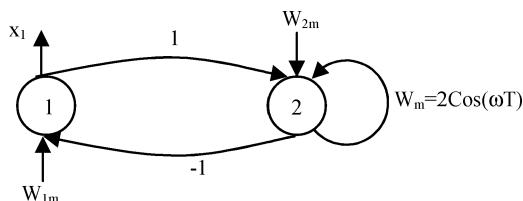

Fig. 2   Single weight oscillator network.

tracking: weights are continuously adapted, or they are adapted in blocks. From a system identification point of view, it is better to adapt in blocks because we will be able to obtain a more precise weight vector at the end of each block. A filter adapted to a given window will be called a local linear model (LLM). However, this approach requires segmentation.

Let us assume that for a given data window N samples long there will exist at most one change in the signal sufficient to require changes in model parameters occurring at time T. The null hypothesis is that the window contains no change; the posited hypothesis is that a change does exist. An example of a memory process similar to this in Sandberg et al.[10] gives us a log-ratio test that can be further simplified for this work. In fact, the input to each of our FIR filters is a sum of unit amplitude, stable sine waves. Therefore, any increase in the error can be univocally assigned to a change in the statistical model of the desired time series. Hence it is not necessary to create two models for the likelihood ratio test. Using maximum-likelihood estimates of model-error variances, a ratio $L(T)$ (referred to as L-values) is defined as

$$2L(T) = (T-1)\ln\left(\sigma_0^2 / \sigma_0'^2\right) + (N-T+1)\ln\left(\sigma_0^2 / \sigma_1^2\right) \quad (13)$$

The errors between the model output and the desired response can be considered to be independently identically distributed. The variance σ_0^2, corresponding to the null hypothesis, is calculated over the entire window. The variances $\sigma_0'^2$ and σ_1^2, corresponding to the posited hypothesis, are calculated from the data before the assumed change time T and after T, respectively.

The algorithm proceeds as follows. First a window of data from the desired response with fixed length N is selected. It is assumed that a database of models exists (with a minimum size of one), one of which correctly models the first portion of the data in the window. The model output is compared with the desired response $d(n)$ in the window, and an error is defined as, $\varepsilon(n) = d(n) - y(n)$, where $y(n)$ is the output of the LLM at sample n for $n = 1, \ldots, N$. The variance σ_0^2 is calculated once using the entire window of data. The other two variances $\sigma_0'^2$ and σ_1^2 are calculated for all times $t = T$ over the entire window as T is moved from one end of the window to the other. The result is a ratio value for each time T in the window. This is repeated for every model in the model database. When the model matches $d(n)$, the ratio L increases rapidly with every sample n. When a model change occurs at time T, the L values for the model originally matched to the signal will drop to very small values, indicating that a change has indeed occurred in the desired signal. The time T where L peaks is identified as the change point.

The final step in the process is to identify the correct model for use after the change, or notify the training algorithm that a new model is required if no current model exists. The postchange model is identified by comparing mean squared error (MSE) values for each model in the database after the change time. If the model's MSE is lower than a preset criteria, it is declared the winner and becomes the new model for the current data from the change point forward. If no model's MSE passes the criteria, the closest model to the minimum is selected as the initial condition for training a new model to match the data window starting at the change point. For training, it must again be assumed that no change in the data occurs over a minimum fixed length of time. Once a new model is trained, it is added to the database, and the process continues with further data windows until the entire data sequence has been evaluated and modeled.

The proper threshold to use when trying to detect a change in a model is sometimes difficult to determine. When a truly matching model exists for the data in the window, the L values become large quite rapidly and are much greater than those calculated for the other models in the database. But when none of the database models are exact, the L values tend to remain small and irregular, especially near either end of the data window. Thresholds that are set too low will result in false alarms.

A switch point that exists very near the beginning of a data window is very difficult to detect because the statistics are not as robust since they are based on a smaller number of samples. A practical
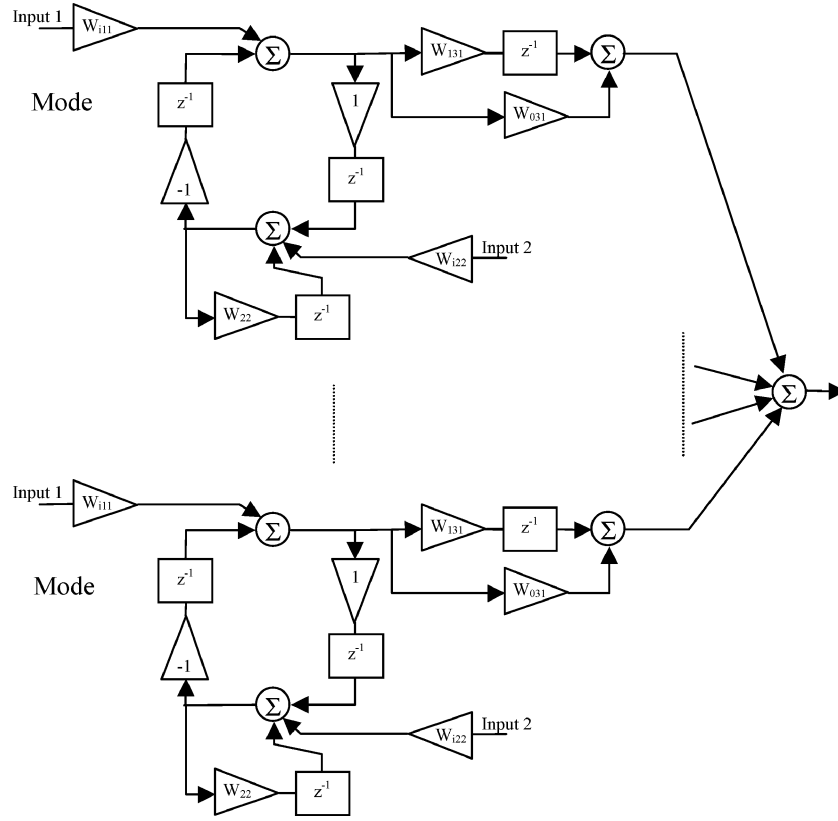
Fig. 3 Complete 16-mode LLM architecture.

Table 1 Flight 8951 free-vibration analysis frequencies by mode

Mode number	Frequency, Hz
1	6.55
2	7.01
3	8.42
4	8.58
5	9.80
6	12.00
7	14.20
8	14.70
9	15.20
10	16.50
11	17.90
12	19.40
13	23.00
14	26.20
15	27.00
16	30.20

minimum of 15 samples was set, where it is assumed that no change has occurred in order for the L values to be stable.

Combining the analysis variance with the network architecture gives us a complete paradigm for modeling a flight-test response as shown in Fig. 4.

Network Initialization and Training Parameters

Prior to training, the frequency-controlling weights in the oscillators were initialized to frequencies identified by structural free-vibration analysis, as shown in Table 1. The FIR zero-order filter weights were initialized to match the free-vibration amplitudes. The first-order weights were set to zero because no phase information is provided by free-vibration analysis.

Training is done using backpropagation, an algorithm whose basis is rooted in the application of the least mean square (LMS). The LMS algorithm makes use of instantaneous estimates of the

autocorrelation and cross correlation of the desired response to estimate the change to a filter weight.[11] Similarly, backpropagation applies a correction to the synaptic weights proportional to the instantaneous error gradient δ_j as

$$\Delta w_{ij}(n) = \eta \delta_j(n) y_i(n), \qquad \delta_j(n) = e_j(n)\phi'_j[v_j(n)] \quad (14)$$

where $e_j(n)$ is the local error at node j, time n, and $\phi'_j[v_j(n)]$ is the first derivative of the activation function at node j, excited by the input v to node j at time n. The learning factor η helps add stability to the solution. The linear nodes used in our networks make this calculation very easy.

Training stability is increased by using resilient backpropagation[12] (RPROP), an algorithm where the size of the weight change varies depending on the slope of the learning curve, the sign of the gradient, and the frequency at which the gradient changes sign. For this work the RPROP adaptation parameters were set to

$$\Delta_{\min} = 10^{-6}, \qquad \Delta_{\max} = 1.0, \qquad \eta_{\min} = 0.5$$

$$\eta_{\max} = 1.2, \qquad \Delta_0 = 10^{-3} \quad (15)$$

Because the network includes recursive components, BPTT was necessary. Updating the weights of recursive networks require that the network be "unfolded." That is, each weight is dependent on the local gradient. Unfolding the network in time is necessary because that gradient is dependent upon not only the current error but each error that has come before it in time.

Network training was accomplished, then, by using BPTT to determine the gradient, then updating the weight using the RPROP rules. The target MSE for training was 10^{-3}. All networks were trained at least three times in order to ensure consistent results. Most LLMs trained within 150 epochs. The actual time required for training each LLM was not tracked. These experiments were accomplished using MATLAB® m-files and as such were not optimized for execution speed. It is understood that a real-time implementation of this paradigm would require high-speed hardware running optimally compiled code because the number of operations is high.
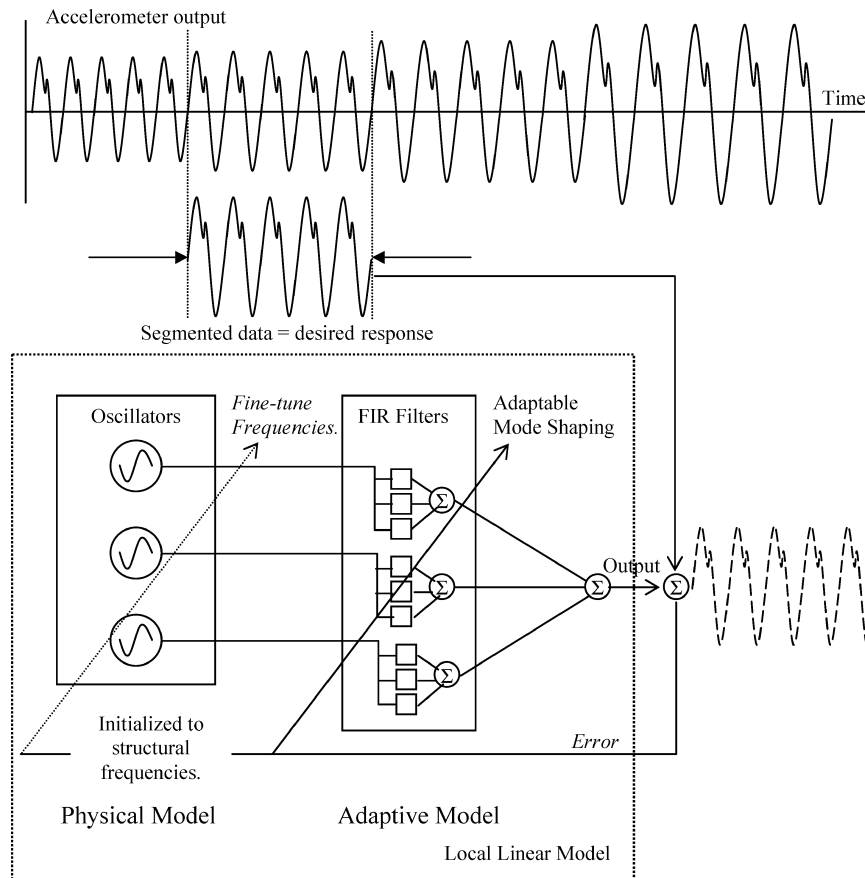
Fig. 4   Hybrid physical model conceptual diagram.

Data Set Description

The process used by the U.S. Air Force[13] for certifying a particular configuration for safety of flight includes engineering analysis and flight testing. Linear flutter analysis seeks to identify whether or not the configuration will result in a flutter or LCO response at a particular set of flight conditions. If analysis indicates that flutter is likely, a properly loaded and instrumented test aircraft is flown under controlled conditions to either confirm the existence of the response or clear the conditions and configuration for flight.

Testing[14] is accomplished by flying the aircraft through a set of flight conditions where active responses are likely to occur. It is not unusual for a single test flight to cover 50 to 75 individual test points. These are usually functions of Mach number, normal acceleration, and altitude. Normal acceleration is an acceleration in a direction normal to the flight path in units of $1\ g \approx 32.2$ ft/s^2 (9.81 m/s^2).

The aircraft is modified with accelerometers and data recorders located at strategic, but practical, points on the aircraft. Our analysis takes its data from the forward wing-tip accelerometer, located approximately 4 in. aft of the wing-tip launcher nose. The measurement range of the accelerometers is $\pm10\ g$ with a nominal sensitivity of 10 mV/g (at 100 Hz) and a noise floor of 0.0012 g (rms). Data are acquired at 216.99 samples per second.

Data for this work are from mission 8951, flown on 28 May 1999 at Eglin Air Force Base, Florida. The aircraft was symmetrically loaded with empty wing-tip launchers, missiles on the underwing station, a 500-lb store on the outer weapon stations, and a fuel tank on the inner weapon stations (Fig. 5). The flight was designed to measure the response of the aircraft under flight conditions ranging from 0.65 to 1.05 Mach, at 2000, 5000, and 10,000 ft, and at g-loadings of $1\ g, 2\ g, \ldots, 6\ g$. The test plan showed about 60 individual test points to be evaluated, organized into 20 lines or groups. Flutter or LCO was expected during the flight, and an abort criterion was set as sustained or increasing vibration at 7 Hz with an amplitude of 3.0 g or higher.

Fig. 5   Aircraft configuration and sensor locations.

Linear Analysis of Flight-Test Configuration

Results of linear flutter analysis using the doublet-lattice method[15,16] of Flight 8951 are shown in Figs. 6 and 7. Predicted damping (Fig. 6) is shown for each mode as a function of knots calibrated airspeed (KCAS). Figure 7 shows the predicted frequencies associated with each mode, also as a function of KCAS. Active modes are those that indicate a definite damping trend, either positive or negative. Modes exhibiting positive damping indicate that some form of flutter can be expected at its frequency, at or near the airspeed where damping becomes positive. Subjectively, modes 1, 2, and 7 could be considered active. Mode 2 is the most active mode and is the only mode predicted to have positive damping beginning at around 350 KCAS. Its frequency at that speed, read as an approximation from the chart in Fig. 7, is 7 Hz. The corresponding free-vibration analysis frequency for that mode is 7.01 Hz. The

Jan 18 2000   f16, wet wing, asym sta 1-LAU-129WT empty   sta 2-LAU-   denegri
                 129/AM AIM-120C-5 sta 3-WEAP-PYL MK-82 LDGP   sta 4-FUEL-P   17
                                  YL WG-TANK 1/4
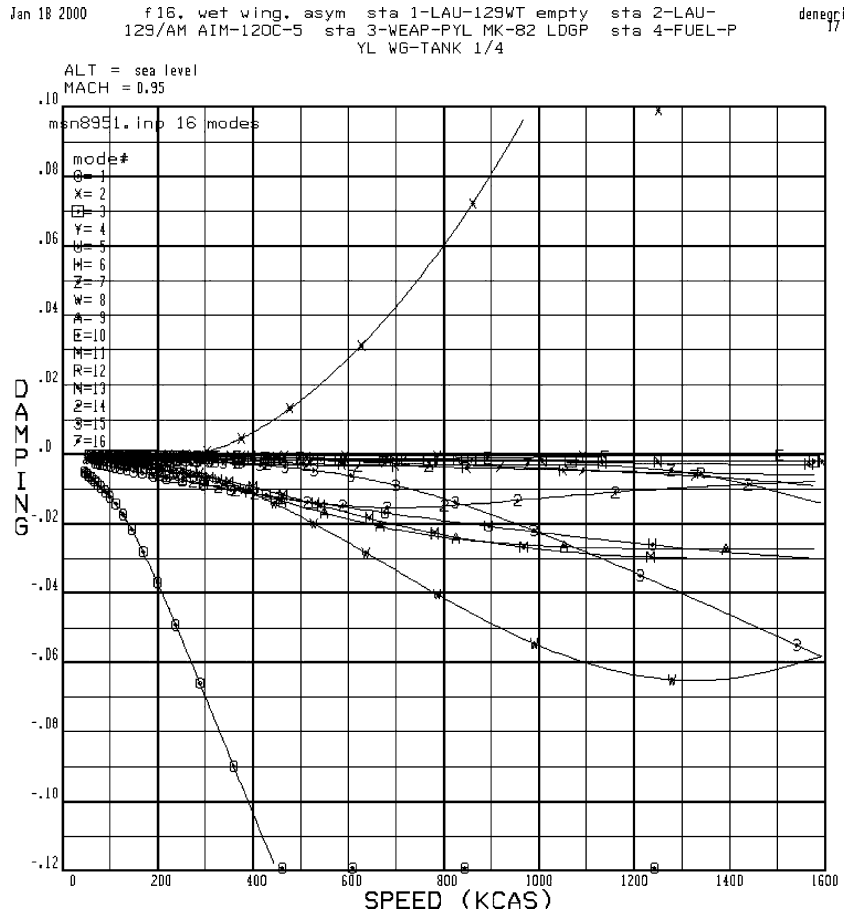
Fig. 6   Linear flutter analysis of test configuration, damping.

free-vibration analysis frequencies for all modes are given in the Table 1.

Flight-Test Signals

Test signals for this work were taken from those measured during flight at test line 17, a signal showing increasing LCO with linearly increasing Mach number. Test line 17 is a portion of flight about 25 s long at 5000 ft mean sea level, in level flight (1 g), with linearly increasing Mach number increasing from 0.78M to 0.9M over about 20 s. Limit-cycle oscillation was observed during this test with increasing amplitude in proportion to increasing Mach number. The flight conditions (test states) are shown in Fig. 8. Altitude in feet, normal acceleration in g, and Mach number are shown as functions of elapsed time for that portion of flight. The normal acceleration plot shows that there was some deviation from the planned value. This is because of normal flight turbulence and is an unfortunate reality of flight testing, resulting in higher noise levels and variations in the data. Mach number was controlled very well, increasing at a nearly linear rate from about 0.78M to 0.9M over 20 s of flight. The Mach number remains constant after 20 s.

Figure 9 shows the forward accelerometer signal for this portion of flight, along with its PSD. As predicted, the response is a form of flutter. In the beginning of the segment, we see that LCO is already present at low but sustained levels. The amplitude of the oscillation increases steadily with increasing Mach number and appears to be flutter at first glance. But a close inspection of the figure shows that as Mach number becomes stable around 0.9, so too do the oscillation amplitudes. The last 5 s are steady, beating limit-cycle oscillations with large maximum amplitudes—around 3.5 g. This was at the abort criteria, but because the response was LCO rather than true flutter the test point was maintained as long as it was deemed safe to do so. The test point was finally terminated 5 s after reaching an airspeed of 0.9M.

As seen in Fig. 9, the primary mode of vibration has a frequency of about 7 Hz as predicted. A second rise in energy is noted at about 14 Hz for this signal. In fact this 14-Hz artifact is noticed in all responses throughout the flight on the forward sensor and is likely to be a second harmonic of the flutter frequency caused by nonlinearities in the system, perhaps structural or fluid. These nonlinearities could be limiting its amplitude, as is generally assumed in all aeroelastic LCO. Linear flutter analysis predicted a mode at 14 Hz, but did not predict it as an active mode. The three modes in the neighborhood of 14 Hz—modes 6, 7, and 8—all were predicted to have negative damping. Consequently, it was a surprise to engineers monitoring the flight to see this mode appear.

The nearly steady flight conditions and the way the LCO response seemed to be in correlation with increasing airspeed made this a good candidate for data to challenge the proposed architecture and sequential neural-network process.

Results

LLM Experiments

Data evaluation begins with the signal from the forward-left wing-tip accelerometer taken during test line 17. The LLM algorithm was given only a single, initial starting network. Figure 10 shows the results of training for the entire signal duration. The MSE value for the trajectory is fairly consistent throughout the entire signal. A maximum MSE value of 10^{-3} was the training goal and was consistently reached in the training sets. As can be seen in the lower window, the overall error and MSE is satisfactory, but is occasionally higher than the desired value because of the errors at switch times. Overall, this family of LLM networks achieved a good synthesized trajectory.

The PSD for the combined network outputs is superimposed on that of the desired response in Fig. 11. Not only is the primary mode correctly represented in the signal by all networks, the PSD of the
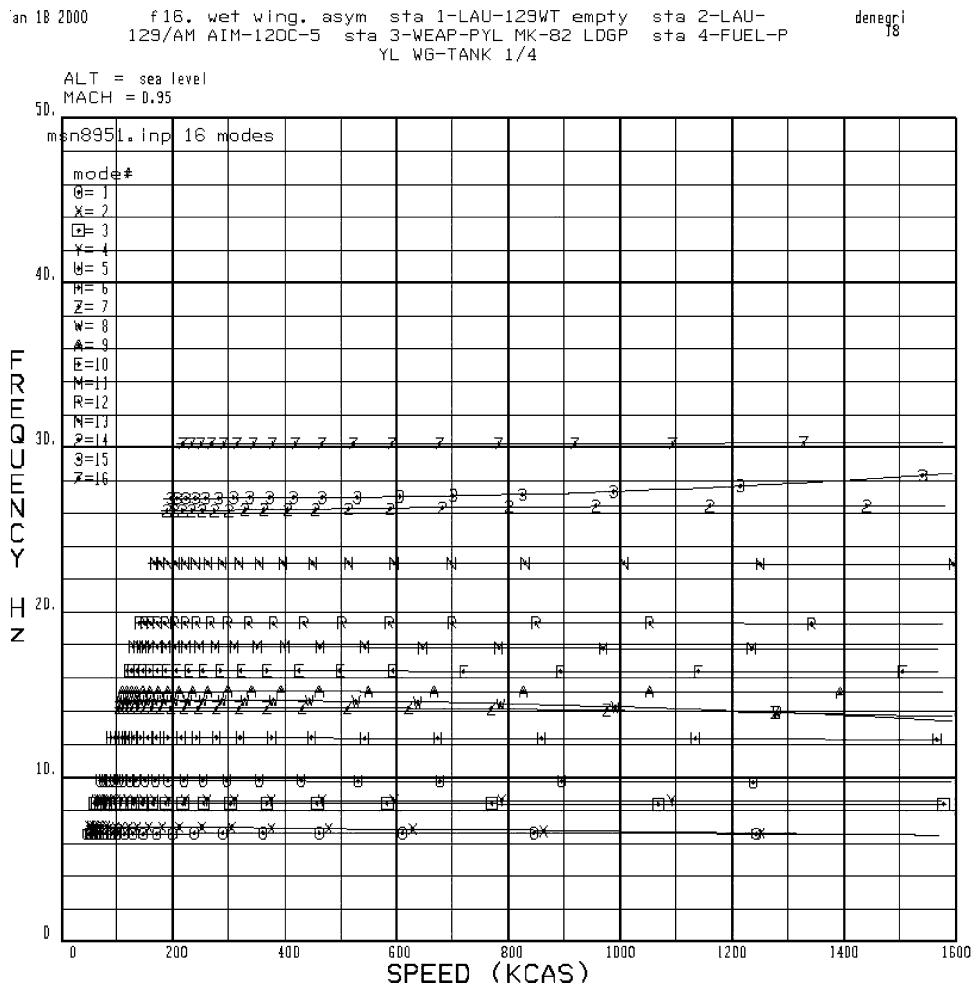
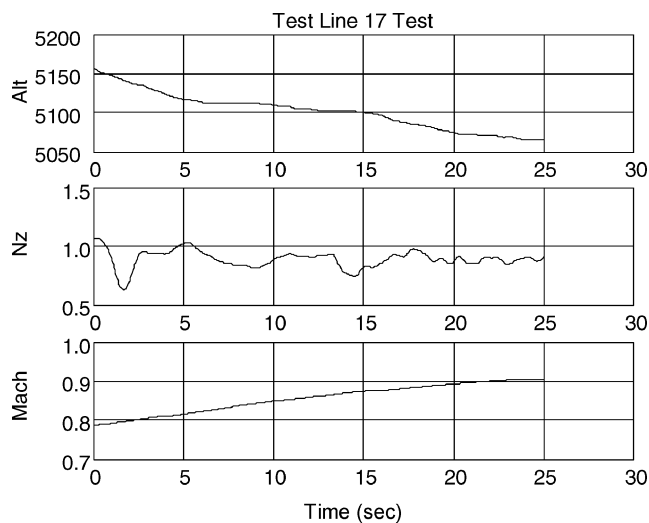**Fig. 7   Linear flutter analysis of test configuration, frequency.**

Fig. 8   Test line 17 test states: altitude, Mach, and normal acceleration.

Fig. 9   Test line 17 forward wing-tip signal and power spectral density.

secondary modes matches the desired response nearly exactly. Slight differences are noted in the energy in the 14-Hz frequency band. The energy is distributed across two of the discrete modal frequencies rather than as a continuous distribution, as would be expected from this type of system.

Because it has been hypothesized that there are only three active modes in this data sequence and that they can be identified through parametric evaluation of the LLMs, it is interesting to see if the output of a three-mode set of LLMs can adequately approximate
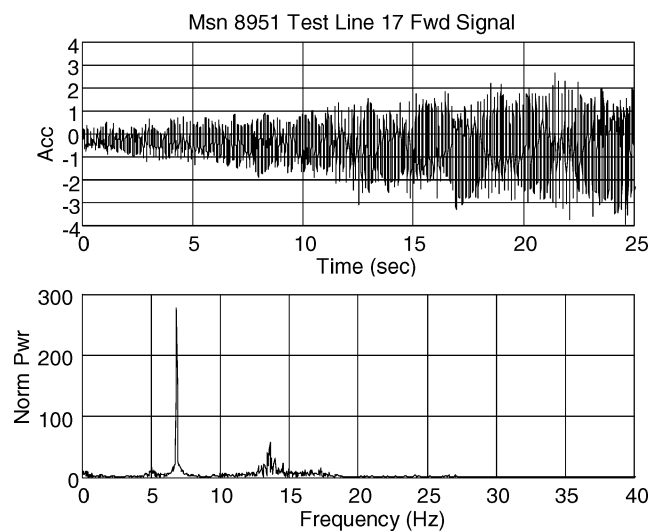
the desired response. The accelerometer data and the output of an adapted three-mode sequence of LLMs are plotted in the first and second windows respectively of Fig. 12.

Several interesting observations can be made. The envelope of the desired response is represented very well by the trimodal approximation, although it tends to exaggerate slightly the more dramatic peaks in the data. These exaggerations are acceptable, however, because from the aspect of safety of flight they represent a realistic but conservative estimation to the signal. Also, the higher frequency noise
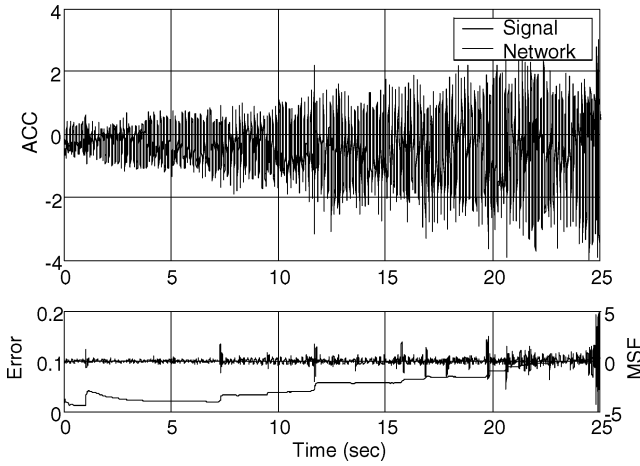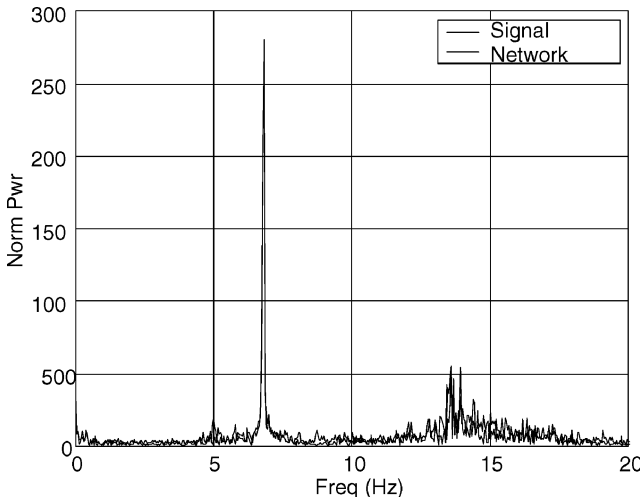
Fig. 10   Test line 17 synthesis.

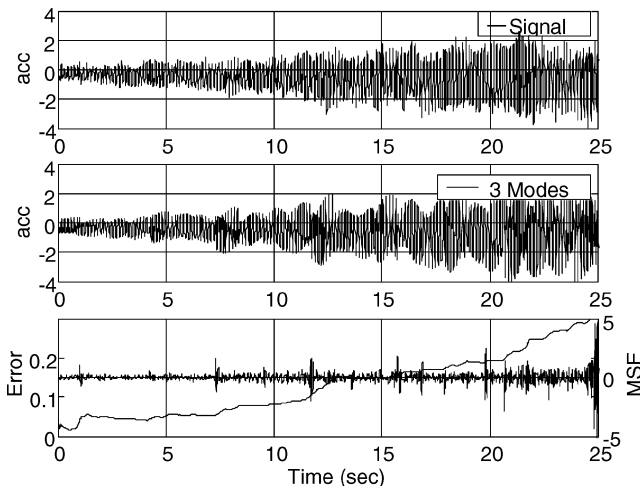Fig. 11   PSD of test line 17 synthesis.

Fig. 12   Data synthesis of test line 17 using LLMs for modes 1, 2, and 7.

has been eliminated while the beating trend in the data is retained. This is of course caused by the limited modal data available for the simulation, but the results provide a clearer picture of the important aspects of the desired response. The MSE stay low throughout the entire trajectory, though higher than that of outputs from the full LLM family of models. The PSD of the trimodal approximation is in Fig. 13. Overall, the trimodal approximation is a reasonable representation of the accelerometer data.
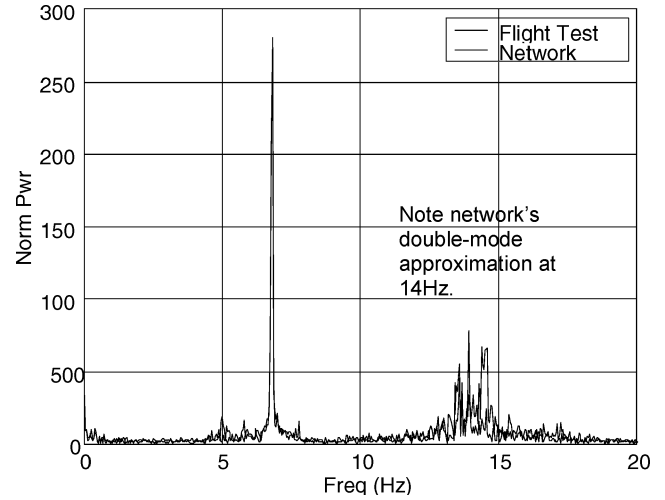

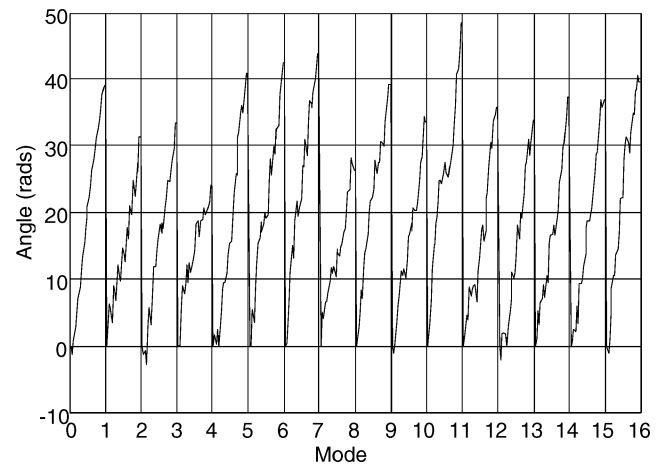Fig. 13   PSD comparison of test line 17 using LLMs for modes 1, 2, and 7.

Fig. 14   Test line 17 LLM FIR modal vector angles by mode.

Parametric Prediction with LLMs

Because the LLM developed here preserves quite a bit of the physical characteristics of the wing flutter, it is appropriate to correlate the change of the filter weights with the physical conditions of the flight. Consider the filter weights as a two-dimensional vector that we will interpret as a phasor that rotates the input modes of the wing. By looking at the real and imaginary parts of weight vectors

$$W_R = \alpha \cos(\phi), \qquad W_I = \alpha \sin(\phi) \qquad (16)$$

we note that the filter weights can be expressed as functions of the filter gain α, the filter phase ϕ, and the modal frequency ω. The one-delay weight is given by

$$W_{21} = \alpha[\sin(\phi)/\sin(-\omega)] \qquad (17)$$

and the zero-delay weight by

$$W_{20} = \alpha \cos(\phi) - W_{21} \cos(-\omega) \qquad (18)$$

expressed compactly as

$$W_{21} = f(\alpha, \phi, \omega), \qquad W_{20} = f(\alpha, \phi, \omega, W_{21}) \qquad (19)$$

If this assumption is verified, a change of an external parameter can be mapped one to one with a system parameter as in any physical model. Therefore, we gain an important insight about the aerodynamic forces in the wing resulting from maneuvers.

The phase angle ϕ is examined for all modes in the entire sequence in Fig. 14. All 400 ϕ values are shown, corresponding to the 25
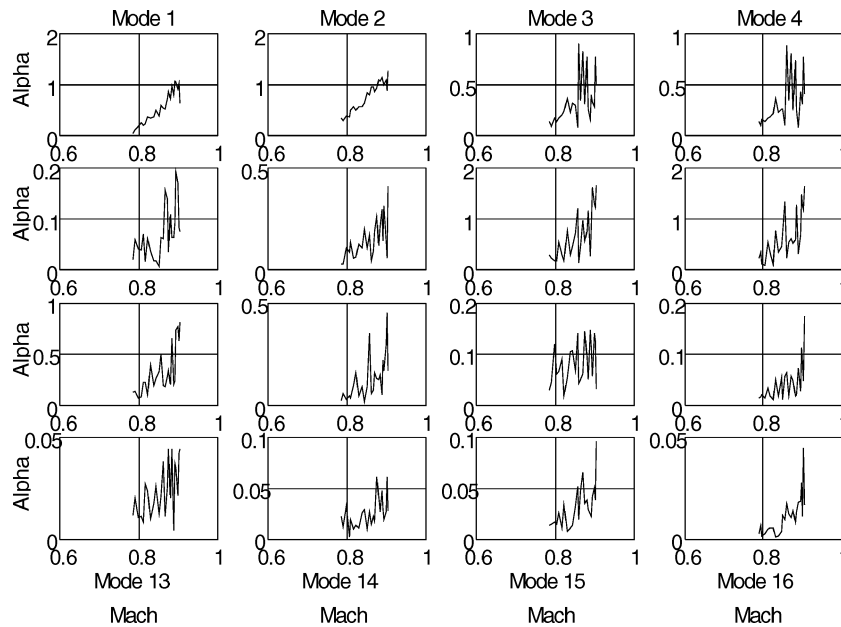
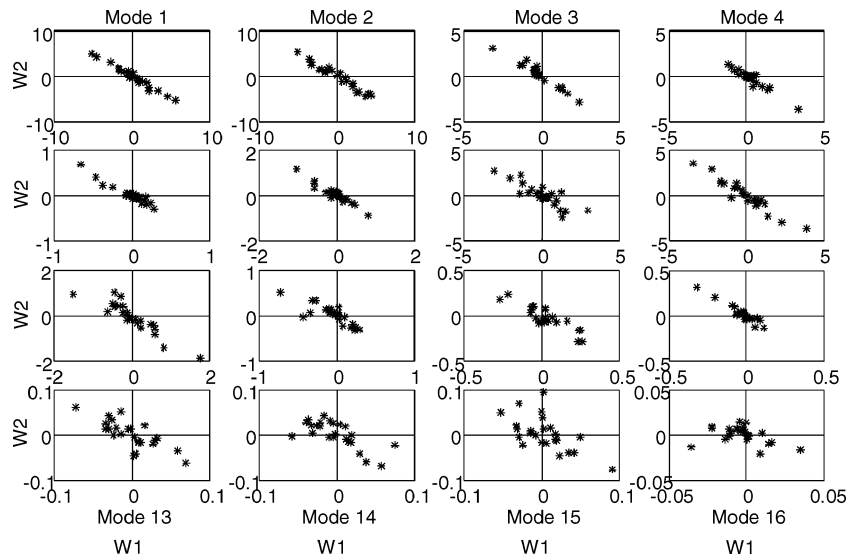Fig. 15   Test line 17 FIR weight vector angles vs Mach.

Fig. 16   Test line 17 FIR zero-order vs first-order weights.

adapted LLMs, each having 16 modal oscillators. The figure shows the angles grouped by mode so that the linear relationship of ϕ for each mode can clearly be observed. For instance, the angles for mode one are the first 25 plotted; those for mode two are 26 through 50, and so forth. The behavior of modes that are active, such as modes 1 and 2, is very linear. Modes that are less active, those do not contribute to the majority of the energy in the signal, have angular responses that are more erratic.

The magnitudes of the FIR filter frequency responses are plotted in Fig. 15 for each mode as a function of Mach number. Again there is a clear linear trend for the active modes. Active modes can be identified by the magnitude of the frequency response compared to other modes in the set. Modes 1 and 2, the modes in the region of the 7-Hz energy band, as well as modes 7 and 8 in the 14-Hz band, have magnitudes much greater than the remaining modes. The magnitude trends of these modes, especially modes 1 and 2, show that they can be approximated linearly. Assuming that this is the case, that these parameters can be modeled as linear functions of Mach, both FIR filter weights for each mode can now be modeled and predicted.

Equation (17) provides a relationship between the zero- and first-order weights. Figure 16 shows these weights plotted one as a function of the other. These weights start out as W_2 set at zero and W_1 set at the structural free-vibration modal amplitude. Because both are modified through adaptation, active modes will show changes in values while inactive modes will move towards zero (or very small values) at the center of each plot. It is clearly seen that when modes are active the relationship between these weights is well behaved. When the modes are not active, the relationship is rather random by comparison. Figure 16 shows that the weights for modes 1 and 2 have much higher values than those corresponding to the remaining modes, as is expected. Modes 3–6 show tight clustering around zero, indicating that for most of the LLMs the adaptive filters are blocking the modal energy. This is also true for modes 10 through 16. The W1 vs W2 plots for modes 7 and 8 show a wider linear spread, again indicating that contribution of these modes is greater. The spread of the weights for mode 9 is wider than those of succeeding modes, but the relationship is less linear, indicative of a more random energy content rather than a truly active mode. Based on observations of
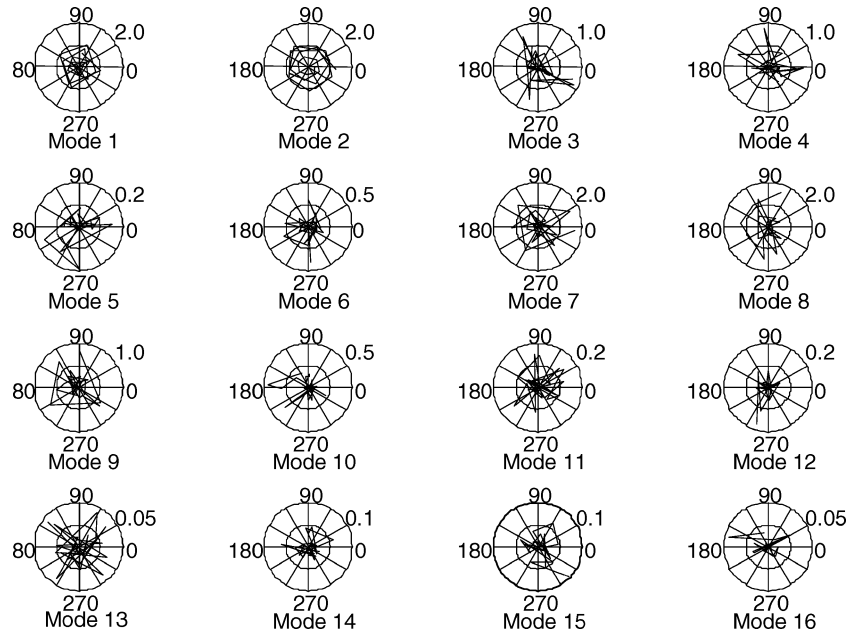
Fig. 17 Test line 17 FIR weight vectors; magnitude and angle.

both Figs. 15 and 16, it can be concluded that the most active modes for test line 17 are 1 and 2 and 7 and 8.

As further evaluation for the selection of active modes and the ability to model LLM parameters, polar plots of the α and ϕ values are shown in Fig. 17. These plots give the clearest indication yet of the active modes by virtue of the total behavior of the FIR weight vectors. As an example, compare the mode 2 and mode 3 plots. The behavior of the mode 2 filter is very circular and well organized, whereas that of mode 3 is random. This circular trend can be seen for modes 1, 2, and 7, whereas the rest of the modes show haphazard, random behavior. It can be concluded from these plots that these three modes are the most active in desired response.

Linearly Fit LLM Parameter Results

Final analysis of the test line 17 data is done by performing a linear fit of the LLM parameters to Mach number as described by Eq. (19). Rather than use detected signal switch times, a linear Mach vector was used to fit the LLM parameters at a Mach resolution of 0.025M. The LLMs were used to construct a data sequence that is compared to the same flight-test accelerometer desired response as before. This comparison is shown in Fig. 18.

The comparison is interesting because it shows both the ability to successfully model the important characteristics of the desired response and the limitations of doing so. The envelope is the most important aspect of the approximation because it is from the envelope that safety of flight for the aircraft configuration is determined. The envelope of the desired response is modeled very well, although influenced by an exaggeration of the beating tendencies in the signal. Throughout the sequence it is seen that any differences between the desired response and the approximation are conservative. That is, the approximations tend to indicate normal accelerations that are slightly higher than those actually encountered, providing a margin of safety. Interesting to note is the high MSE and errors shown. Because the synthesized signal is a result of models changing at regular intervals rather than at switch times, the desired response and model outputs tend to drift out of phase.

The PSD for the linear-fit LLMs is shown in Fig. 19. It is clear from a graphical comparison of the frequency contents of the desired response and its approximation that energy is distributed more evenly between the three active modes in the LLMs. All modes are represented, but the linear approximation of LLM parameters and the more even distribution of LLMs as a function of Mach number have effectively removed all but the three active modes of oscillation.
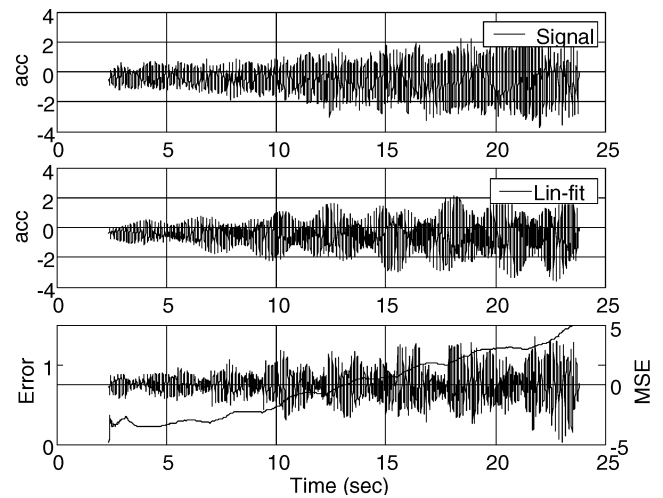

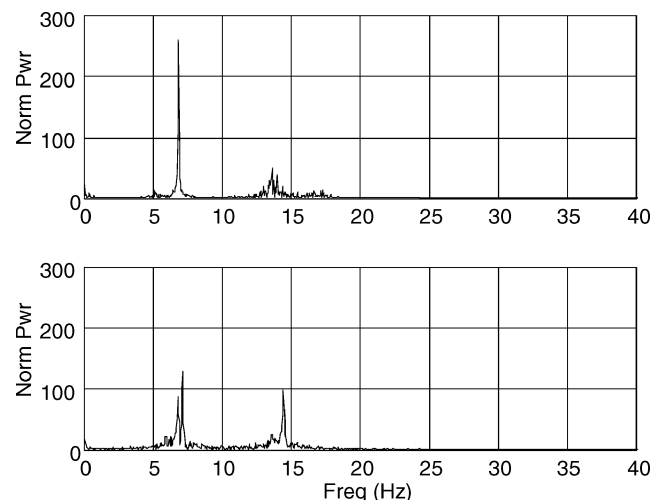Fig. 18 Data synthesis of test line 17 using linear-approximated LLMS.

Fig. 19 PSD comparison of data synthesis using linear-approximated LLMs.

Table 2   Test line 17 modes 1 to 8 adapted frequencies

Mode 1	Mode 2	Mode 3	Mode 4	Mode 5	Mode 6	Mode 7	Mode 8
			Free-vibration frequency, Hz				
6.5460	7.0091	8.4156	8.5785	9.8033	12.3948	14.2342	14.6988
			Adapted frequencies, Hz				
6.5453	7.0085	8.4161	8.5780	9.8038	12.3951	14.2338	14.6985
6.5463	7.0079	8.4173	8.5775	9.8033	12.3936	14.2341	14.6985
6.5481	7.0072	8.4162	8.5780	9.8046	12.3907	14.2339	14.6987
6.5446	7.0064	8.4153	8.5786	9.8077	12.3893	14.2337	14.6988
6.5516	7.0079	8.4140	8.5796	9.8092	12.3880	14.2338	14.7002
6.5445	7.0098	8.4165	8.5798	9.8031	12.3975	14.2336	14.6986
6.5440	7.0078	8.4144	8.5807	9.8066	12.3883	14.2344	14.6990
6.5437	7.0098	8.4152	8.5761	9.8050	12.3933	14.2340	14.6986
6.5463	7.0089	8.4174	8.5793	9.8030	12.3952	14.2339	14.6977
6.5444	7.0087	8.4143	8.5784	9.8142	12.3928	14.2325	14.6990
6.5438	7.0096	8.4162	8.5782	9.8080	12.3882	14.2343	14.6986
6.5477	7.0062	8.4192	8.5764	9.8038	12.3904	14.2356	14.6991
6.5450	7.0087	8.4126	8.5757	9.8020	12.3895	14.2326	14.6996
6.5470	7.0080	8.4136	8.5767	9.8020	12.3929	14.2336	14.6981
6.5477	7.0064	8.4106	8.5722	9.7974	12.3897	14.2348	14.6987
6.5458	7.0076	8.4241	8.5823	9.8000	12.3910	14.2337	14.6988
6.5449	7.0084	8.4202	8.5825	9.8020	12.3948	14.2339	14.6981
6.5451	7.0088	8.4216	8.5825	9.8007	12.3932	14.2338	14.6984
6.5466	7.0078	8.4145	8.5774	9.7997	12.3928	14.2336	14.6976
6.5441	7.0093	8.4146	8.5766	9.8039	12.3925	14.2336	14.6985
6.5443	7.0092	8.4185	8.5755	9.8089	12.3897	14.2338	14.6989
6.5459	7.0084	8.4162	8.5795	9.8094	12.3905	14.2343	14.7000
6.5441	7.0099	8.4154	8.5772	9.8101	12.3934	14.2338	14.6987
6.5407	7.0102	8.4211	8.5809	9.8127	12.3892	14.2352	14.7004
6.5407	7.0102	8.4211	8.5809	9.8127	12.3892	14.2352	14.7004

The 7-Hz primary mode is now made up of the energy contributions of two close modes.

Results of test line 17 LLM data synthesis and analysis of the adapted LLM parameters show the potential of the LLM paradigm for use online for predicting wing responses based on previous and current model parameters. Further, it is seen that examination of LLM components can provide insight into the critical makeup of the modal behavior of LCO responses.

Frequency Adaptation Analysis

There is some expectation that the modal frequencies of the LCO oscillations can shift over the course of flight. The network architecture and its adapted parameters allow a unique evaluation of these frequencies. The frequency controlling weights of the linear oscillators provide a direct calculation for each mode. The results of a frequency analysis are interesting. They show that although the relationship between LCO and flight test states can be nonlinear, the frequencies stay nearly constant throughout the flight. Table 2 shows the frequencies for the test line 17 data, modes 1 through 8. Modes 9 through 16 are shown in Table 3.

The frequencies for the active modes vary only slightly from LLM to LLM, an indication that for practical purposes the modal frequencies are remaining consistent over the course of the flight, regardless of changing dynamics that are encountered. This makes intuitive sense because frequency is heavily influenced by the structural characteristics of the wing's mass and stiffness.

Model Generalization Evaluation

As an exercise to evaluate the generalization of the LLMs, the first model for test line 17, associated with a normal acceleration of 1 g and Mach number of 0.78, flown at 5000 ft, is applied to anther set of data, the first portion of test line 3 lasting about 1.8 s. This test was also at 1 g, with a Mach number of 0.8, and at an altitude of 10,000 ft. The comparison is shown in Fig. 20. The LLM from test line 17 does not fit the data exactly, but does provide a close enough approximation to give very reasonable estimates of the LCO that is encountered. Most important to note is that the bias is not necessary. The bias noticed during test line 17 is not present in this data set. Therefore it must be ignored when applying the model.
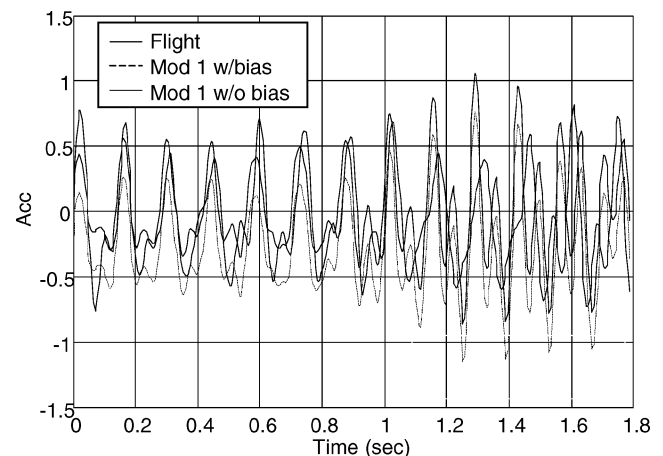
Fig. 20   Data synthesis of test line 3 segment 1, using test line 17 model 1.

There are two critical differences between the conditions under which these tests were conducted that can help to explain this. Intuitively it would be expected that the second LLM from test line 17 would be the better match because its Mach number is also 0.8. But the calibrated airspeeds for the two are not the same because of the altitude differences. The calibrated airspeed value associated with the test line 17 (0.78M) model is about 453 KCAS, whereas that of test line 3 segment 1 (0.8M) is 450 KCAS. Because calibrated airspeed accounts for the differences in altitude between the two sets, it makes sense that the first LLM of test line 17 is a closer match to test line 3 segment 1.

The second critical difference lies in the manner of testing. Test line 17 is a test flown in level flight with a consistently increasing airspeed. It can be argued that other variables of flight will be changing as well. Most significant is angle of attack. As airspeed increases with all other variables remaining constant, angle of attack decreases in a similar, linear fashion. Thus a bias in the direction of the LCO responses can be expected. Contrast this to the test line 3 conditions. The goal of that test was to maintain constant normal

Table 3   Test line 17 modes 9 to 16 adapted frequencies

Mode 9	Mode 10	Mode 11	Mode 12	Mode 13	Mode 14	Mode 15	Mode 16
Free-vibration frequency, Hz							
15.2114	16.4557	17.8903	19.3882	22.9792	26.1874	26.9629	30.2262
Adapted frequencies, Hz							
15.2111	16.4560	17.8900	19.3884	22.9794	26.1876	26.9630	30.2264
15.2110	16.4550	17.8912	19.3886	22.9799	26.1866	26.9624	30.2255
15.2112	16.4552	17.8921	19.3876	22.9804	26.1867	26.9614	30.2251
15.2122	16.4564	17.8919	19.3904	22.9802	26.1886	26.9615	30.2247
15.2119	16.4556	17.8917	19.384	22.9798	26.1879	26.9623	30.2257
15.2097	16.4548	17.8913	19.3879	22.9799	26.1864	26.9639	30.2258
15.2127	16.4565	17.8913	19.3914	22.9798	26.1901	26.9628	30.2243
15.2105	16.4565	17.8907	19.3891	22.9802	26.1877	26.9613	30.2268
15.2114	16.4566	17.8916	19.3892	22.9792	26.1862	26.9628	30.2256
15.2123	16.4567	17.8926	19.3838	22.9810	26.1912	26.9627	30.2258
15.2131	16.4570	17.8914	19.3909	22.9786	26.1896	26.9626	30.2238
15.2122	16.4573	17.8852	19.3900	22.9810	26.1836	26.9649	30.2299
15.2093	16.4549	17.8913	19.3938	22.9780	26.1901	26.9642	30.2216
15.2102	16.4529	17.8913	19.3909	22.9798	26.1906	26.9670	30.2216
15.2096	16.4533	17.8914	19.3926	22.9798	26.1910	26.9687	30.2226
15.2084	16.4547	17.8926	19.3893	22.9815	26.1883	26.9566	30.2292
15.2105	16.4569	17.8921	19.3833	22.9832	26.1910	26.9592	30.2296
15.2101	16.4573	17.8920	19.3819	22.9819	26.1904	26.9598	30.2296
15.2088	16.4533	17.8912	19.3914	22.9785	26.1901	26.9679	30.2208
15.2101	16.4552	17.8903	19.3910	22.9778	26.1931	26.9672	30.2204
15.2120	16.4572	17.8957	19.3921	22.9796	26.1943	26.9627	30.2260
15.2126	16.4567	17.8929	19.3893	22.9774	26.1900	26.9651	30.2252
15.2111	16.4554	17.8920	19.3913	22.9814	26.1848	26.9614	30.2271
15.2143	16.4601	17.8995	19.3924	22.9770	26.1891	26.9609	30.2294
15.2143	16.4601	17.8995	19.3924	22.9770	26.1891	26.9609	30.2294

acceleration, airspeed, and altitude for a period of time. It is expected, then, that angle of attack will also remain constant and the bias in LCO would not be as noticeable, certainly not comparable to that of test line 17.

The best fit occurs during the first second of the data. There appears to be a phase shift in the desired response at about 1.25 s. The linear model is not able to adjust to follow. The envelope of the desired response is well represented by the LLM, though, and, as with prior models, would provide an accurate yet conservative estimate of the expected LCO response. The overall MSE for the model is 0.204 for the LLM when the bias is included in the LLM and 0.168 without applying the bias.

Summary

The experiments using the paradigm and model architecture to synthesize an extended flutter response show that a sequence of component-based local linear models consisting of quasi-fixed IIR pole-placing subnetworks driving a set of adaptive FIR filters can be used to synthesize the IIR LCO response of the wing under dynamic aeroelastic forces and can provide descriptions of all modes contained in the model, effectively decomposing the flight-test signal into its modal representation.

The computational burden is eased by a number of factors. The LLMs, once initialized, are used as initial conditions for the next LLM in the sequence. This means that the adaptation process does not have to search the entire solution space every time a model is trained. It simply has to fine tune the previous model. When networks are originally initialized, or when dynamics rapidly and severely change, it was shown that linear flutter analysis results provide a good starting point from which to search for a final solution.

The paradigm provides several indicators that identify the active modes in the accelerometer signals. As a test of the accuracy of these, signals were synthesized using only LLMs corresponding to those modes identified as active. There were three modes identified as active. Models for these were used, and it was shown that the signals could be synthesized very well. The trimodal LLMs resulted in a signal simulation that could be used to determine safe flight conditions because in all cases the error between modeled and actual responses was conservative. These determinations are subjective, but data synthesis was deemed successful if the question of whether

or not the synthesized signal could be used by engineers observing its behavior to determine approximate LCO conditions and safety of flight.

Confirmation of Existing Methods

This work relied a great deal upon current practices for the prediction and evaluation of LCO and flutter in thin wing aircraft. Linear flutter analysis is still the backbone of the industry. It is simply too expensive to use other methods for practical flutter analysis at this time. Because traditional methods are so heavily depended upon, it is good to be able to provide insight into their accuracy whenever the opportunity presents itself. The fundamental basis for linear flutter analysis is that the solutions of its equations can be solved by assuming that flutter, and by extension LCO, can be considered a phenomenon that can be locally modeled as a sum of sinusoids. Based on the success of the modularly designed LLMs to synthesize the flight-test data, those LLMs themselves representing a sum of sinusoids, it is reasonable to conclude that the assumption made by linear flutter methods that flutter can be modeled as such is correct. It was shown that this assumption holds for steady-state LCO responses at least for fixed-state variables, that is, altitude, Mach number, normal acceleration, and aircraft configuration.

Further, structural free-vibration analysis provides modal frequencies that are in close proximity of the actual values. Based on the analysis of the modal content of the data used, it can be concluded that linear flutter theory's approximation of these frequencies is valid. It was shown that when the linear oscillator networks were initialized to these values very little adjustment was required.

Further Analyses

Several topics for further research and experimentation are suggested from the results presented in this work. The experiments conducted showed the viability for the LLM paradigm for data synthesis, as well as opening new points of view on the construct of the LCO phenomenon. It follows that further evaluation of the use of parametric modeling and approximation can help to answer some of the questions that deal with the signals and their characteristics.

At the present time linear flutter analysis considers all forces on the wing to be linear combinations of structural or aerodynamic

influences. Consider the information already obtained from the experiments in this work. The linear relationship of model parameters to Mach tends to reinforce the original assumptions. The implication is that by modeling more data under a wider variety of flight conditions and then analyzing the relationships between the adapted LLMs as functions of independent variables can provide insight into the nature of the forces involved, perhaps sufficiently to further define the current empirical models.

There were some preliminary experiments conducted during the course of this work with other data sets, particularly the test line 3 set used in the parametric comparison described, that were not specifically reported on because of insufficient data from which to draw conclusions. There was a trend noted, however, that is interesting to mention. The test line 3 data were taken at discrete g loadings from 1 to 4 g. Networks were trained for portions of each flight condition, although the data were partitioned manually rather than by an analysis of variances into five separate, disconnected segments. In general, a third-order fit was required to describe the network parameters. Results of simulations made with the approximated network parameters did not match the flight data as well as the experiments in this work were able to, and it was concluded that although a fit was likely more data were needed to accurately describe the parametric relationships. Therefore a higher-order trend in the FIR filter weights as a function of g loading is implied, but cannot be shown conclusively without more data and analysis.

The most practical use for this work is as an online model change detection and prediction scheme for flight safety determination. Currently a flight must be monitored at all times by engineers watching for tell-tale signs of the onset of unsafe flight conditions. By using the paradigms described, a prediction method can be set up providing the expected response ahead of the flight conditions, either to provide a warning or to ensure continued safe testing. It was shown that under carefully controlled conditions this is possible and viable. Because all testing must be done under extremely controlled conditions, the utility of this paradigm is immediate. All tests are flown with data telemetry equipment that allows real-time data processing on the ground, at high speed, during the flight. The methods presented here can be implemented and used with the technology currently available. Fail-safes can be built in to ensure against continued flight in unsafe conditions when the predictions are uncertain, when data transmission fails, etc.

Conclusions

The usefulness of this paradigm as an online tool was clearly demonstrated. Noting that the signal from test line 17 was very long reinforces this conclusion. The system was able to continuously model a changing wing response without fail for 25 s. The results of experiments using that data, exploiting the concept of building a continuous series of sequential local linear models to model a signal, can easily be interpreted to imply that the system could be used to model signals for an indefinite period of time, provided that the hardware used can keep up with the calculations required.

It was shown that for portions of flight whose conditions are changing on Mach number the network parameters could be modeled as linear functions of Mach. The ability to predict model parameters as a function of flight-test state provides a capability to the network for extrapolating wing response beyond the current conditions. This was demonstrated during the test line 17 experiments with very good success. The finer details of the signal were not

present, but the overall envelope of the signal as well represented throughout the entire 25-s test duration.

Finally, an evaluation was made intended to demonstrate the ability of a model adapted with one set of data to generalize to a similar but unknown data set. The first adapted model for the test line 17 data was applied to the first data sequence from test line 3. These two tests had similar flight conditions, but differed in Mach number and altitude. The results showed that the model's output, although not an exact match, was fully sufficient to predict the trend in a limit-cycle oscillation response from which a safety of flight judgment could be made. Errors with this experiment were on the conservative side, as they were with all experimental results.

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]Cunningham, A. M., Jr., "Semi-Empirical Unsteady Aerodynamic 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.
[3]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.
[4]Denegri, C. M., Jr., Dubben, J. A., and Maxwell, D. L., "In-Flight Wing Deformation Characteristics During Limit Cycle Oscillations," *Journal of Aircraft*, Vol. 42, No. 2, 2005, pp. 500–508.
[5]Fung, Y. C., *An Introduction to the Theory of Aeroelasticity*, Dover, Nineola, NY, 1955, p. 192.
[6]Cunningham, H. J., Batina, J. T., and Bennett, R. M., "Modern Wing Flutter Analysis by Computational Fluid Dynamics Method," *Journal of Aircraft*, Vol. 25, No. 10, 1988, pp. 962–968.
[7]Tang, L., Bartels, R. E., Chen, P. C., and Liu, D. D., "Simulation of Transonic Limit Cycle Oscillations Using a CFD Time-Marching Method," AIAA Paper 2001-1290, April 2001.
[8]Johnson, M. R., and Principe, J. C., "A Two-Processing Element Adaptable Linear Oscillating Recurrent System with Single-Weight Plasticity," Inst. of Electrical and Electronics Engineers, Paper 0-7803-7898-9/03, July 2003.
[9]Principe, J. C., Euliano, N. R., and Lefebvre, W. C., *Neural and Adaptive Systems, Fundamentals of Through Simulations*, Wiley, New York, 2000, p. 208.
[10]Sandberg, I. W., Lo, J. T., Fancourt, C. L., Principe, J. C., Katagiri, S., and Haykin, S., *Nonlinear Dynamical Systems*, Wiley, New York, 2001, p. 153.
[11]Haykin, S., *Neural Networks A Comprehensive Foundation*, Macmillan, New York, 1994, p. 144.
[12]Riedmiller, M., and Braun, H., "A Direct Adaptive Method for Faster Backpropagation Learning: The RPROP Algorithm," *Proceedings of the IEEE International Conference on Neural Networks*, edited by H. Ruspini, ICNN, San Francisco, 1993, pp. 586–591.
[13]Dyess, W. W., Jr., and Stokes, D., "The Role of Digital Models and Simulations in the Air Force SEEK EAGLE Office," *2001 Advanced Simulation Technologies Conference*, Manuscript Number M172, Society for Computer Simulation International, Seattle, WA, April 2001.
[14]Dreyer, C. A., and Shoch, D. L., "F-16 Flutter Testing at Eglin Air Force Base," AIAA Paper 86-9819, April 1986.
[15]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.
[16]Desmarais, R. N., and Bennett, R. M., "An Automated Procedure for Computing Flutter Eigenvalues," *Journal of Aircraft*, Vol. 11, No. 2, 1974, pp. 75–78.