

Using Distributed Fiber-optic Strain Sensing to Estimate Generalized Modal Coordinates from Flight-test Data

Jeffrey Ouellette,¹ Matthew Boucher,² and Peter Suh³

NASA Armstrong Flight Research Center, Edwards, California, 93523, U.S.A.

A method for estimating the generalized modal coordinates of an aircraft during flight has been developed. The Fiber-optic Sensing System (FOSS) offers an efficient and cost-effective method of measuring the strain at thousands of points along the wings. The estimation of modal coordinates was implemented as a two-step process. First, a maximum likelihood method is used to estimate the statistical properties of the sensors and generalized modal coordinates. Second, the strain mode shapes from the finite element model are used along with the statistical properties from the first step to estimate the generalized modal coordinates over time. Using simulated data from the X-56A Multi-Utility Technology Testbed (MUTT), different methods of modal coordinate estimation were compared to demonstrate the benefits and weaknesses of each. These were compared against the exact solution and an ordinary least squares (a more traditional method) estimate. Modal coordinate estimation methods were then applied to flight-test data from the X-56A aircraft to show that the method continues to work well with actual test data. The new estimation method provides insights unavailable from more classical approaches.

I. Nomenclature

AMLG	antisymmetric main landing gear mode
AW1B	antisymmetric wing 1st bending mode
AW1T	antisymmetric wing 1st torsion mode
AW2B	antisymmetric wing 2nd bending mode
Boom Lat	air data boom lateral mode
Boom Vert	air data boom vertical mode
CFD	computational fluid dynamics
corr	correlation matrix of data
cov	covariance matrix of data
$E(\cdot)$	expectation
$\mathfrak{F}$	factor scores
FA	factor analysis
FEM	finite element model
FFT	fast Fourier transform
FOSS	Fiber-optic Sensing System
$\mathbf{I}$	identity matrix
IID	identically and independently distributed
$\mathbf{J}$	measurement noise weighting matrix
$\mathfrak{K}$	step size scaling
KKT	Karush-Kuhn-Tucker conditions
$\mathfrak{L}$	factor loadings
$\mathcal{L}(\cdot)$	negative log-likelihood

¹Aerospace Engineer, Controls and Dynamics Branch, P.O. Box 273/Mailstop 4840D, Edwards, California, 93523, AIAA Member.

²Aerospace Engineer, Controls and Dynamics Branch, P.O. Box 273/Mailstop 4840D, Edwards, California, 93523.

³Aerospace Engineer, Controls and Dynamics Branch, P.O. Box 273/Mailstop 4840D, Edwards, California, 93523.

M	mass matrix
MAC	modal assurance criterion
MUTT	Multi-Utility Technology Testbed
n	number of samples
NG	nose gear stroke mode
OLS	ordinary least squares
P	generalized coordinate correlation matrix
PCA	principle components analysis
POD	proper orthogonal decomposition
R^2	coefficient of determination
RMS	root mean square
S	sample covariance matrix of outputs
S_T	singular value matrix
SFA	symmetric wing fore-aft mode
Sim	simulation
SMLG	symmetric main landing gear mode
SVD	singular value decomposition
SWL	symmetric winglet bending mode
SW1B	symmetric wing 1st bending mode
SW1T	symmetric wing 1st torsion mode
T	factor rotation matrix
$\text{tr}(\cdot)$	trace of a matrix
U_T	singular value left eigenvectors
V	covariance of measurement noise
$\mathbf{v}$	measurement noise or specific factors
V_T	singular value right eigenvectors
$W_p(\cdot)$	Wishart random variable
Z	matrix of outputs
$\mathbf{z}$	outputs
$\bar{\mathbf{z}}$	sample mean

Subscripts

k	index of sample
m	number of modes or factors
p	number of measurements

Superscripts

$*$	intermediate estimate of the parameter
ij	the element in the i -th row and j -th column of the matrix
ols	ordinary least squares estimate
reg	regression estimate
wls	Bartlett's (weighted least squares) estimate

Symbols

∂	partial derivative
$\boldsymbol{\eta}$	generalized coordinate
$\boldsymbol{\kappa}$	mode shape to factor scaling matrix
$\bar{\boldsymbol{\kappa}}$	mode shape to rotated factor scaling matrix
$\boldsymbol{\Lambda}$	scaling of mode shapes
$\boldsymbol{\Phi}$	mode shape
$\bar{\boldsymbol{\Phi}}$	rotated/target mode shapes
$\boldsymbol{\Sigma}$	covariance matrix of outputs
$\boldsymbol{\mu}$	mean of output
$\mathcal{V}$	eigenvectors

Conventions

$\oslash$	element-wise (Hadamard) division
$\arg \min$	minimizing argument of the expression
diag	diagonal matrix
$\ \cdot\ _2$	2-norm
$\dagger$	pseudoinverse
$\hat{(\cdot)}$	estimate of parameter
$\succ$	positive definite matrix
$\triangleq$	equal by definition

II. Introduction

THERE is a trend in aircraft development to create lighter and more fuel efficient structures, which is resulting in increasing flexibility and lower structural margins. In an effort to increase the high altitude performance of the RB-57H Canberra (General Dynamics Corporation) (Reston, Virginia) airplane, shown in Fig. 1. The size of the wing was increased; however, the larger wing was more flexible which caused body freedom flutter where the short-period mode was destabilized by this increased flexibility [1].



Fig. 1 The RB-57H Canberra airplane (General Dynamics Corporation) (Reston, Virginia) (Photo credit: NASA/JSC).

The effect of increased wing flexibility is especially apparent in unconventional vehicle configurations such as high aspect ratio flying wings like the Helios (AeroVironment Inc.) (Arlington, VA) unmanned aerial vehicle, shown in Fig. 2. Large structural deformations from a wind gust significantly reduced the stability of the phugoid mode and caused the vehicle to break up in flight [2]. Better measurement of the state of the structure could have mitigated the mishap by allowing the control system to suppress or otherwise correct for the effect of the large deformations.



Fig. 2 The Helios (AeroVironment Inc.) (Arlington, Virginia) unmanned aerial vehicle in flight (Photo credit: NASA/Carla Thomas).

Although the increased flexibility presents new challenges in aircraft design, there are also many developments in new sensing technologies such as the Fiber-optic Sensing System (FOSS), which can enable thousands of spatially distributed structural strain measurements. The highly distributed nature of the FOSS offers many unique opportunities to improve awareness of the structural deformations, but robust methods are needed to transform the data into a useful signal.

The standard approach in structural dynamics, is modal analysis where the deformations are separated into a spatial mode shape, Φ , and a temporal generalized modal coordinate, η_k . The separation of variables transforms the partial differential equations into separate solvable ordinary differential equations [3]. The objective of the present work is to estimate these generalized modal coordinates. The ability to better estimate the structural state will enable many new technologies in structural control, e.g., flutter suppression, gust load alleviation, maneuver load alleviation, and new methods of system identification for these flexible vehicles.

These generalized modal coordinates are a latent variable, meaning they cannot be directly measured or observed. The process of estimating these latent generalized modal coordinates from test data is often called modal filtering. The use of a modal filter offers a method of reducing the thousands of measurements to a smaller set of generalized coordinates that can be used more practically by a control system. Previous modal filtering methods [4, 5] have typically utilized an ordinary least squares (OLS) method for estimating the generalized coordinates. However, noise can cause the OLS approach to significantly overestimate the response of some modes. As a result, the OLS is only able to estimate coordinates of the most dominant modes, and may be incapable of estimating the generalized coordinates of other secondary modes that are contributing to the dynamics. The truncation of these secondary modes can also lead to spillover of the modes where these truncated modes corrupt the estimates of the primary modes [4]. Grauer [4, 5] uses accelerometer measurements and a Kalman filter to correct the measurements for the noise. Yet, the OLS estimation feeding into the Kalman filter will still be susceptible to spillover of truncated modes and errors in the mode shapes from the finite element model.

Unlike previous modal filtering, the present methodology first generates an estimate of the statistical properties of the measurement noise and the generalized modal coordinates independent of the estimate of the time histories of the noise and modal coordinates. The independent estimate of these statistical properties allows a more informed estimation of the generalized modal coordinates. Additionally, the present work is better able to handle errors in the finite element model (FEM) mode shapes by estimating independent mode shapes based on patterns in the data. A rotation is then defined which correlates these data modes to the FEM mode shapes. While previous modal filtering studies have used theoretical models [4] or wind tunnel models with a simplified structure [6, 5], the present method has been applied to flight-test data of the X-56A (Lockheed Martin) (Bethesda, Maryland) Multi-Utility Technology Testbed (MUTT), shown in Fig. 3. The X-56A MUTT had a realistic monocoque structure with spars, ribs, and skin. These structural elements, like the skin and rib attachments, cause local discontinuities in the strain that contradict Saint Venant's principle [7]. This makes the elements difficult or impossible to capture in FEM models of a complete vehicle used in aeroelastic analysis. With a realistic monocoque structure, testing structural sensing technologies like FOSS and

the analysis of the data with a realistic monocoque structure is important for assessing the robustness to such model errors.



Fig. 3 The X-56A MUTT in flight (Photo credit: NASA/Jim Ross).

Section III., “Background on Factor and Modal Analyses,” provides background on the origin of these modal and factor analysis techniques described to provide more clarity on the assumptions leading to their creation, their similarities, and their differences. Section IV., “Definition of Factor / Modal Model,” defines the mathematical model for the estimation problem. Section V., “Fitting Methodology,” describes the optimization algorithm for determining the parameter values. In section VI., “X-56A Results,” the proposed method has been applied to flight-test data from X-56A flex-wing flights.

III. Background on Factor and Modal Analyses

The common factor model used in factor analysis is a general concept where measured variables are assumed to be a combination of the influence of a smaller number of unobservable latent variables. In the early 20th century, the statisticians Karl Pearson and Charles Spearman conceived a model that was mathematically similar to modal analysis while trying to define and study human intelligence [8].

The factor model is equivalent to the way in which the deformations of an aircraft are assumed to be the linear combination of mode shapes. The primary difference in these two models is that in modal analysis, the unobservable generalized coordinates, η_k , are expected to obey some ordinary differential equation. Also, in modal analysis the modes are to separate the spatial component and therefore define a shape - known as a mode shape. Factor analysis (FA) is more general, and the factor loadings can be any common construct of the outputs. Despite the differences in the origin of the methods, because factor analysis has developed largely independent of the history of modal analysis, means that there are many lessons that can be adapted to benefit the application of modal analysis to test data. It is assumed most readers will be more familiar with the nomenclature of modal analysis, so in general, this formatting will be used throughout the present work. The similarity of the terminology between modal and factor analysis is shown in Table 1.

The use of modal analysis in engineering dates back to the work of Joseph Fourier in the 19th century and predates the concept of the common factor model. The idea of modal analysis stems from the attempts to solve the partial differential equations that are used to describe the dynamics of a structure in space and time. As these equations were an attempt to solve a theoretical mathematics model, there is no measurement noise, $\mathbf{v}_k \equiv \mathbf{0}$. As long as people have been analyzing the structural dynamics, there has been a need to validate those structural dynamics. However, many of these early methods of modal analysis were very challenging and required the estimation of each mode one at a time [9]. The advent of the fast Fourier transform (FFT) [10] opened up new possibilities in the field of modal testing. The FFT allows the identification of multiple modes from a single test.

Methods of estimating these latent generalized modal coordinates have been developed for ground vibration testing.

However, these tests usually require a very large number of accelerometers fixed to the vehicle, and it is difficult and expensive to instrument a flight vehicle with this many sensors. Furthermore, when implemented in flight, these sensors return complex valued aeroelastic modes. These modes may be similar to the ground modes but will change with airspeed, and without further analysis or the use of a model, it is impossible to relate them back to the correlated ground test modes. The use of the FOSS does improve the number of sensors in flight, but the data sampling rate is limited and returns strain measurements. The generation of shape or deflections requires additional processing of the data. The assumptions of the additional processing for deflections introduces further errors.

Table 1 Similarity of terms from factor and modal analyses.

Modal analysis		Factor analysis	
	Mode		Common factor
Φ	Mode shape	$\mathcal{L}$	Factor loadings
η	Generalized coordinate	$\mathfrak{F}$	Factor scores
$\mathbf{v}$	Measurement noise	$\mathbf{v}$	Specific factors

A. Proper Orthogonal Decomposition (Principal Component Analysis)

A closely related technique to factor analysis is principal components analysis (PCA), also frequently referred to as proper orthogonal decomposition (POD) among engineers [11, 8]. In many cases, the PCA/POD is a reasonable approximation to the factor analysis solution. The PCA/POD solution is identical to the first iteration of a principal factor solution. However, factor analysis and PCA/POD are intended to address fundamentally different questions. Because of the similarities and the popularity of PCA/POD, it is worth discussing the differences and why factor analysis was selected over PCA/POD.

The primary objective of PCA/POD is to simplify the data by identifying patterns in the data. These patterns could be used as factor loadings in a factor analysis. However, the PCA/POD is not indented to characterize a latent variable like a generalized modal coordinate. Also, there is no explicit handling of measurement noise. The PCA/POD methods have gained popularity in engineering to reduce the order of computational fluid dynamics (CFD) simulations. Because CFD does provide the data at all grid points without measurement noise, the PCA/POD is an appropriate method. The PCA/POD method is very appealing because it is based on a lone singular value decomposition (SVD); this makes it a numerically stable and efficient method of identifying patterns.

In contrast, in factor analysis we are trying to identify the underlying relationship, but we still have patterns in the outputs as a result of the factor loadings. However, these outputs do not represent all the data. The FOSS system provides a lot of measurements of strain, but certainly, not every measurement. There are variations in the strain between the measurements reported by the FOSS. Furthermore, because the goal is to work with experimental data, there can be significant noise. For many of the strain measurements, especially near the X-56A wingtip, the signal to noise is less than one.

IV. Definition of Factor / Modal Model

In modal analysis, the measured outputs, $\mathbf{z}_k$, are a linear combination of the generalized modal coordinates, η_k , plus the measurement noise, $\mathbf{v}_k$, in Eq. (1). All biases, including sensor bias and strain from static deformations, are grouped together into a single mean parameter, $\bar{\mathbf{z}}$. The goal of the present work is to determine the value of the generalized modal coordinates, η_k , of Eq. (1). Estimates of these generalized modal coordinates, the strain mode shapes, Φ , and the magnitude of the measurement noise, $\mathbf{v}_k$, also need to be determined. The strain mode shapes, Φ , are taken from the FEM.

The common factor model, Eq. (2), appears to be identical to the modal model in Eq. (1). The only difference between the factor model and the modal model is in the definition of the scaling and normalization.

$$\underset{(p \times 1)}{\mathbf{z}_k} = \underset{(p \times m)}{\Phi} \underset{(m \times 1)}{\eta_k} + \underset{(p \times 1)}{\mathbf{v}_k} + \underset{(p \times 1)}{\bar{\mathbf{z}}} \quad (1)$$

$$\underset{(p \times 1)}{\mathbf{z}_k} = \underset{(p \times m)}{\mathbf{L}} \underset{(m \times 1)}{\mathbf{f}} + \underset{(p \times 1)}{\mathbf{v}_k} + \underset{(p \times 1)}{\bar{\mathbf{z}}} \quad (2)$$

The normalization of the factor model differs from the modal model because the factor scores, $\mathbf{f}$, are normalized, Eq. (3a). In contrast to factor analysis where the factor scores are normalized, in modal analysis the mode shapes, Φ , are normalized, Eq. (3b). The mode shapes are typically normalized with respect to the mass matrix, $\mathbf{M}$. Furthermore, in factor analysis often the factors are assumed to be orthogonal with respect to themselves.

$$\text{diag}(\text{Cov}(\mathbf{f})) = \mathbf{I} \quad (3a)$$

$$\Phi^T \mathbf{M} \Phi = \mathbf{I} \quad (3b)$$

To account for the difference in normalization, a scaling matrix, κ , is defined, Eq. (4), to give the factor loadings, $\mathbf{L}$.

$$\mathbf{L} \triangleq \Phi \kappa \quad (4)$$

The standard orthogonal normal modes in structural analysis are orthogonal with respect to the mass and stiffness matrices. However, the time histories of the modes, η , are not generally orthogonal. The mode time histories will only be orthogonal for some special cases such as a string or membrane. To accommodate these oblique modes, the oblique factor model is needed, and the covariance of the factor scores is not assumed to be orthogonal. In the oblique factor model, the matrix, $\mathbf{P}$, is the correlation matrix, Eq. (5).

$$\mathbf{P} \triangleq \text{Cov}(\mathbf{f}) \equiv \text{Corr}(\eta) \quad (5)$$

The modal coordinates are assumed to be of zero mean, Eq. (6). In general, the generalized coordinates do not have a zero mean. However, by subtracting the mean of each output, is assumed to provide a sufficient approximation to satisfy this assumption.

$$E(\eta) = \underset{(m \times 1)}{\mathbf{0}} \quad (6)$$

The noise on each output is assumed to be independent Gaussian white noise, Eqs. (7a) and (7b), with mean of zero and covariance $\mathbf{V}$.

$$E(\mathbf{v}) = \underset{(p \times 1)}{\mathbf{0}} \quad (7a)$$

$$E(\mathbf{v}\mathbf{v}^T) = \underset{(p \times p)}{\mathbf{V}} = \begin{bmatrix} \nu_1 & 0 & \cdots & 0 \\ 0 & \nu_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \nu_p \end{bmatrix} \quad (7b)$$

Additionally, the measurement noise is assumed to be independent of the modal coordinate, Eq. (8). The condition of Eq. (8) is not actually true for the FOSS. The physics of the laser used in the FOSS causes the noise to be dependent on the loading applied; however, the assumption is reasonable and effective.

$$E(\mathbf{v}\eta^T) = \underset{(p \times m)}{\mathbf{0}} \quad (8)$$

In factor analysis, rather than looking at the modal coordinates, η_k , we consider first the covariance matrix of the output, Σ , defined in Eq. (9).

$$\underset{(p \times p)}{\Sigma} = \underset{(p \times m)}{(\Phi \kappa)} \underset{(m \times m)}{\mathbf{P}} \underset{(m \times p)}{(\Phi \kappa)^T} + \underset{(p \times p)}{\mathbf{V}} \quad (9)$$

Under the above assumptions, the estimate of the covariance of the output, $\hat{\Sigma}$, is a Wishart random variable, Eq. (10).

$$\hat{\Sigma} \sim W_p(\mathbf{S}, n) \quad (10)$$

The covariance of the outputs can also be approximated by the sample covariance, $\mathbf{S}$ in Eq. (11).

$$\mathbf{S} \triangleq \frac{1}{n} \sum_{k=1}^n \mathbf{z}_k \mathbf{z}_k^T \quad (11)$$

The parameter n in Eq. (10) is the number of samples.

Rather than directly optimizing the probability density of the Wishart distribution of Eq. (10), the negative log-likelihood function, $\mathcal{L}$, of the parameters of the Wishart distribution, Eq. (12), is used for the optimization [8, 12, 13, 14].

$$\mathcal{L}(\Sigma|\mathbf{S}) = \frac{n}{2} (\log |\Sigma| + \text{tr}(\Sigma^{-1}\mathbf{S})) + \text{constant} \quad (12)$$

V. Fitting Methodology

The fundamental problem of this work is determining the value of the parameters in the factor model, Eq. (1). In this section, first, the method of minimizing the negative log-likelihood function, Eq. (12), is described; second, the estimated mode shapes are correlated back to the FEM mode shapes by a rotation matrix; and third, these parameter estimates are used to estimate the generalized modal coordinates. The maximum likelihood method, where the parameters are selected as the values which are most likely to produce the data measured, was selected. The maximum likelihood solution is the solution that minimizes Eq. (12). The maximum likelihood solution is beneficial because it has a strong basis in statistical methods, which provide a strong basis and a well-established set of tools for assessing the model. However, the maximum likelihood model is generally the most complex cost function to fit.

In the process of determining the parameter values, some values may be known from previous analysis. For example, the mode shapes can be calculated from the FEM. Furthermore, it is possible to make assumptions on the parameters that can reduce the number of parameters that need to be fit. Reducing the number of parameters can make the fitting process easier and reduce uncertainty. As a result, it is desirable for the parameters to be either a constant (known) or a varying value and determined by the estimation process.

A. Maximum Likelihood Factor Analysis

The fitting process of the maximum likelihood model generally requires the gradients of the negative log-likelihood, Eq. (12). These gradients for the standard factor model are derived by Jöreskog [15] and Lawley and Maxwell [12]. The gradients are expanded to include the scaled factor model with κ . The gradients shown are in Eqs. (13a)-(13d).

$$\frac{\partial \mathcal{L}}{\partial \Phi} = -n \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa \mathbf{P} \kappa \quad (13a)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{P}} = -\frac{n}{2} \kappa \Phi^T \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa \quad (13b)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{K}} = & -\frac{n}{2} \left[\Phi^T \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa \mathbf{P} \right] \\ & -\frac{n}{2} \left[\Phi^T \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa \mathbf{P} \right]^T \end{aligned} \quad (13c)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{V}} = -\frac{n}{2} \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \quad (13d)$$

The maximum likelihood solution are the parameters where these gradients go to zero.

1. Quasi-Newton Method

The first method considered was a direct use of a commercial quasi-Newton method solver. The quasi-Newton solver was appealing because it appeared to be a directly applicable solver and is an algorithm with well-established numerical robustness. The quasi-Newton method was extremely slow to converge to a solution, and the numerical accuracy was very poor. As a result, the quasi-Newton method was not used to generate the current results. The application of the quasi-Newton method is still described because it may be more useful in some future work.

The first problem that arose with the quasi-Newton method was the numerical condition on the matrices of the factor model, Eqs (14a)-(14e).

$$\mathbf{P} \succ 0 \quad (14a)$$

$$\mathbf{V} \succ 0 \quad (14b)$$

$$\text{diag } \mathbf{P} = \mathbf{1} \quad (14c)$$

$$\mathbf{V} \text{ is diagonal} \quad (14d)$$

$$\boldsymbol{\kappa} \text{ is diagonal} \quad (14e)$$

These numerical conditions are reflected as constraints on the optimization. The conditions of Eqs. (14d) and (14e) are very easy to implement by constructing the matrix from a vector of parameters. The condition on the diagonal of $\mathbf{P}$, Eq. (14c), was straightforward to implement as an equality constraint in the optimization. It is possible to leave the diagonal terms of the $\mathbf{P}$ matrix out of the parametrization. However, this makes the construction of the $\mathbf{P}$ matrix significantly more difficult, so the equality constraint was used instead.

The most challenging constraint was the positive definition condition, Eq. (14a). When looking at the positive definiteness of a matrix, the eigenvalues are often used. The solution of the eigenvalues is an efficient numerical algorithm, and positive eigenvalues are easy to verify. The evaluation of gradients of the eigenvalues, required for the nonlinear optimization methods, is extremely challenging. The ordering of the eigenvalues can change in the finite difference methods, making it impossible to calculate the gradients. Furthermore, the occurrence of repeated eigenvalues makes attempts to sort the eigenvalues impossible.

To avoid the challenges with eigenvalues, Sylvester's criterion was used to constrain the positive definiteness instead. Sylvester's criterion requires the determinate of the leading principal minors of the matrix to be positive. Sylvester's criterion is a necessary and sufficient condition for the positive definiteness of a matrix. Sylvester's criterion is much more straightforward to implement. However, in this case, the determinants of the minors were not well conditioned. The different minors could have significantly different magnitudes. The quasi-Newton methods struggled to handle these constraints with very different magnitudes.

The second issue that arose with the quasi-Newton optimization was in the numerical accuracy of the solution. Each of the terms in the factor model can have substantially different magnitudes. The terms of $\mathbf{P}$ are generally on the order of -1 to 1 . However, the terms of $\boldsymbol{\kappa}$ were often on the order of 10^3 or larger. These differences in magnitude made most convergence criteria in the optimizer ineffective. Because very small changes in $\mathbf{P}$ would have a large effect on the solution, the optimization would converge more effectively for these parameters, while $\boldsymbol{\kappa}$ would not be well converged. When the optimizer would consider the solution converged, the $\boldsymbol{\kappa}$ matrix would have much larger numerical errors than the $\mathbf{P}$ matrix. These problems could have been improved by defining better convergence criteria but would have required significant rework of the optimizer and reduced the "off-the-shelf" benefit of using an existing quasi-Newton optimizer.

The final issue with the quasi-Newton method was the extremely slow convergence. The slow convergence was due in part to the numerical problems. Very small changes in the parameters could make very large changes in the cost function. The sensitivity of the cost function made accurate evaluations of the gradients difficult, so the optimizer was forced to make very small steps when close to the solution. Furthermore, a tighter tolerance was needed to achieve a satisfactory level of convergence. Ultimately, the convergence was much too slow, more than ten times as long as the block coordinate descent method below.

2. Block Coordinate Descent

To solve the maximum likelihood factor analysis, a method based on the approach of Lawley [12] called block coordinate descent was used. This method was initially rejected because it would not converge consistently. The method would often oscillate between solutions, but once a fix was found to the oscillating solutions problem, the block coordinate descent became a much more effective solution method. As a result, this block coordinate descent method was used to generate all of the results in this paper.

The Karush–Kuhn–Tucker (KKT) conditions, Eqs (15a)-(15c), provide necessary conditions for the maximum likelihood solution of the factor model.

$$\frac{\partial \mathcal{L}}{\partial \boldsymbol{\Phi}} = 0 \quad (15a)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{P}} = 0 \quad (15b)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{V}} = 0 \quad (15c)$$

First, the mode shapes, Φ , are considered. Although the FEM does provide an estimate of the mode shapes, there are errors in the mode shapes. To handle these errors, it can be beneficial to generate a new set of shapes which are then related back to the original FEM mode shapes in section V.B., “Rotations of Shapes.” To estimate a new set of shapes, in Eq. (16a) the gradient from Eq. (13a) is applied to the KKT condition of Eq. (15a). The $\kappa \mathbf{P} \kappa$ is non-singular and can be removed, Eq. (16b). Finally, the addition is separated on the left and right sides of the equation to give a simpler necessary condition for the solution of the mode shapes, Eq. (16c).

$$\Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa \mathbf{P} \kappa = 0 \quad (16a)$$

$$(\Sigma^{-1} - \Sigma^{-1} \mathbf{S} \Sigma^{-1}) \Phi = 0 \quad (16b)$$

$$\implies \Phi = \mathbf{S} \Sigma^{-1} \Phi \quad (16c)$$

The mode shapes do still remain a function of itself and the unknown factor model covariance, Σ . This is a problem that will be addressed later.

The next term is the generalized modal coordinate correlation matrix, $\mathbf{P}$. An intermediate matrix $\mathbf{J}$, Eq. (17), is defined that will be used to simplify the expressions for the modal coordinate correlation matrix.

$$\mathbf{J} \triangleq \Phi^T \mathbf{V}^{-1} \Phi \quad (17)$$

In Eq. (18a) the gradient from Eq. (13b) is applied to the KKT condition of Eq. (15b). The Sherman-Morrison-Woodbury (matrix inversion) Lemma [16] is used to expand the inverse of the factor model covariance, Eq. (18b). Non-singular square-matrices are removed in Eq. (18c) without the loss of generality of the solution. The factor covariance matrix is expanded by the definition in Eq. (18d). The matrix $\mathbf{J}$, Eq. (17), is used to simplify the expanded expression in Eq. (18e). The terms without the $\mathbf{P}$ matrix are moved to the right side of the expression, Eq. (18f). Finally, the inverse of the $\mathbf{J}$ is multiplied on the left and right to isolate the $\mathbf{K}$ and $\mathbf{P}$ matrices on the left, Eq. (18g).

$$\kappa \Phi^T \Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} \Phi \kappa = 0 \quad (18a)$$

$$(\mathbf{I} + \mathbf{J} \mathbf{P})^{-1} \kappa \Phi^T \mathbf{V}^{-1} (\Sigma - \mathbf{S}) \mathbf{V}^{-1} \Phi \kappa (\mathbf{I} + \mathbf{J} \mathbf{P})^{-1} = 0 \quad (18b)$$

$$\implies \Phi^T \mathbf{V}^{-1} (\Sigma - \mathbf{S}) \mathbf{V}^{-1} \Phi = 0 \quad (18c)$$

$$\implies \Phi^T \mathbf{V}^{-1} (\Phi \kappa \mathbf{P} \kappa \Phi^T + \mathbf{V} - \mathbf{S}) \mathbf{V}^{-1} \Phi = 0 \quad (18d)$$

$$\mathbf{J} \kappa \mathbf{P} \kappa \mathbf{J} + \mathbf{J} - \Phi^T \mathbf{V}^{-1} \mathbf{S} \mathbf{V}^{-1} \Phi = 0 \quad (18e)$$

$$\mathbf{J} \kappa \mathbf{P} \kappa \mathbf{J} = \Phi^T \mathbf{V}^{-1} \mathbf{S} \mathbf{V}^{-1} \Phi - \mathbf{J} \quad (18f)$$

$$\kappa \mathbf{P} \kappa = \mathbf{J}^{-1} \Phi^T \mathbf{V}^{-1} \mathbf{S} \mathbf{V}^{-1} \Phi \mathbf{J}^{-1} - \mathbf{J}^{-1} \quad (18g)$$

The conditions on the κ and $\mathbf{P}$ matrices do make this solution unique. The $\mathbf{P}$ matrix is defined as the correlation matrix of the generalized modal coordinates, the diagonals are all one. The κ matrix is therefore uniquely determined by the square root of the diagonal of the right side of Eq. (18g). The κ matrix is then used to scale Eq. (18g) to calculate the correlation matrix, $\mathbf{P}$.

The final term is the output noise measurement covariance, $\mathbf{V}$. In Eq. (19a) the gradient from Eq. (13d) is applied to the KKT condition of Eq. (15c). The non-singular matrices are removed, and the expression is expanded with the factor model definition, Eq. (19b). Finally, the $\mathbf{V}$ matrix is separated on the left side of the expression, Eq. (19c).

$$\Sigma^{-1} (\Sigma - \mathbf{S}) \Sigma^{-1} = 0 \quad (19a)$$

$$\Phi \kappa \mathbf{P} \kappa \Phi^T + \mathbf{V} - \mathbf{S} = 0 \quad (19b)$$

$$\mathbf{V} = \mathbf{S} - \Phi \kappa \mathbf{P} \kappa \Phi^T \quad (19c)$$

Necessary conditions have been defined for each term of the maximum likelihood factor model; however, these terms are each a function of the remaining terms. This makes an analytical solution impossible. Instead, an iterative coordinate descent method has been used. The coordinate descent optimization is used by minimizing the cost function one

parameter, or a block of parameters, at a time. The parameters are grouped by the matrices in the factor model and solved in the order of Eqs. (20a)-(20e).

$$\Sigma_{(k)} = \Phi_{(k)} \kappa_{(k)} \mathbf{P}_{(k)} \kappa_{(k)} \Phi_{(k)}^T + \mathbf{V}_{(k)} \quad (20a)$$

$$\Phi_{(k+1)}^* = \mathbf{S} \Sigma_{(k)}^{-1} \Phi_{(k)} \quad (20b)$$

$$\mathbf{J}_{(k)}^* = \kappa_{(k)} \Phi_{(k+1)}^T \mathbf{V}_{(k)}^{-1} \Phi_{(k+1)} \kappa_{(k)} \quad (20c)$$

$$\mathbf{P}_{(k+1)}^* = \kappa_{(k)} \mathbf{J}_{(k)}^{-1} \left(\kappa_{(k)} \Phi_{(k+1)}^T \mathbf{V}_{(k)}^{-1} \mathbf{S} \mathbf{V}_{(k)}^{-1} \Phi_{(k+1)} \kappa_{(k)} - \mathbf{J}_{(k)} \right) \mathbf{J}_{(k)}^{-1} \kappa_{(k)} \quad (20d)$$

$$\text{diag } \mathbf{V}_{(k+1)}^* = \text{diag} \left(\mathbf{S} - \Phi_{(k+1)} \kappa_{(k+1)} \mathbf{P}_{(k+1)} \kappa_{(k+1)} \Phi_{(k+1)}^T \right) \quad (20e)$$

The oscillatory/zig-zag behavior initially observed is a common problem in first-order optimization methods like steepest decent. The oscillations can occur if the solution in one coordinate overshoots the minimum. The solution path then tends to jump wildly about the parameter space. By reducing the step size, the solution in one coordinate block will not overshoot the global solution.

In the event that a matrix has a constant parameter, or constraint, the solution is not updated. The update scheme is summarized in Eqs. (21a)-(21d).

$$\Phi_{(k+1)}^{ij} = \begin{cases} \Phi_{(k)}^{ij} & \text{if constant} \\ \frac{1}{\kappa+1} \left(\Phi_{(k+1)}^{*ij} + \Re \Phi_{(k)}^{ij} \right) & \text{otherwise} \end{cases} \quad (21a)$$

$$\mathbf{P}_{(k+1)}^{ij} = \begin{cases} 1 & i = j \\ \mathbf{P}_{(k)}^{ij} & \text{if constant} \\ \frac{1}{\kappa+1} \left(\frac{\mathbf{P}_{(k+1)}^{*ij}}{\sqrt{\mathbf{P}_{(k+1)}^{*ii} \mathbf{P}_{(k+1)}^{*jj}}} + \Re \mathbf{P}_{(k)}^{ij} \left(\kappa_{(k)}^{ij} \right)^2 \right) & \text{otherwise} \end{cases} \quad (21b)$$

$$\kappa_{(k+1)}^{ij} = \begin{cases} 0 & i \neq j \\ \kappa_{(k)}^{ij} & \text{if constant} \\ \frac{1}{\kappa+1} \left(\sqrt{\mathbf{P}_{(k+1)}^{*ii} \mathbf{P}_{(k+1)}^{*jj}} + \Re \kappa_{(k)}^{ij} \right) & \text{otherwise} \end{cases} \quad (21c)$$

$$\mathbf{V}_{(k+1)}^{ij} = \begin{cases} 0 & i \neq j \\ \mathbf{V}_{(k)}^{ij} & \text{if constant} \\ \frac{1}{\kappa+1} \left(\mathbf{V}_{(k+1)}^{*ij} + \Re \mathbf{V}_{(k)}^{ij} \right) & \text{otherwise} \end{cases} \quad (21d)$$

Typically, the coordinate descent methods are considered to be relatively slow algorithms, only providing linear convergence. The reduction of the step size further slows the overall rate of convergence. Because the solution expressions, Eqs. (20a)-(20e), are of a much lower computational cost than the gradients or Hessian, the time for an individual iteration is much smaller than a quasi-Newton method. The faster iterations of the block coordinate descent results in a lower overall solution time (about ten-times faster) than a higher order method, like quasi-Newton.

3. Initial Estimates

Because all the current fitting methods are iterative, a reasonable initial guess is required. The method, based on the approach of Johnson and Wichern [8], was easy to implement and generally gave good results. Johnson and Wichern's approach was expanded to include the scaling matrix, to permit constant mode shapes, and to permit oblique factors.

The initial output noise, Eq. (22), is based on the magnitude of the sample covariance, scaled is based on the ration of the number of modes and the number of outputs.

$$\mathbf{V}_{(0)}^{ij} = \frac{1 - \frac{m}{p}}{\text{diag}(\mathbf{S}^{-1})} \quad (22)$$

When the mode shapes are varying, the modes are assumed orthogonal, so the generalized modal coordinate correlation matrix, $\mathbf{P}_{(0)}$, is defined to be identity. When constant modes are used, the initial correlation matrix, Eq. (23), is calculated with the pseudo-inverse.

$$\mathbf{P}_{(0)} = \begin{cases} \left(\Phi^T \mathbf{V}_{(0)} \right)^\dagger \mathbf{S} \left(\mathbf{V}_{(0)} \Phi \right)^\dagger & \text{constant mode shapes} \\ \text{Identity} & \text{otherwise} \end{cases} \quad (23)$$

The initial guess of the mode shapes is shown in Eq. (24).

$$\Phi_{(0)} = \begin{cases} \Phi & \text{constant mode shapes} \\ \mathbf{V} & \text{otherwise} \end{cases} \quad (24)$$

Where $\mathbf{V}$ is taken from the eigenvalue decomposition, eigenvalues Λ and eigenvectors $\mathbf{V}$, of the weighted sample matrix in Eq. (25). This is calculated using the eigenvalue decomposition. In the decomposition, the eigenvalues are sorted from the largest to smallest values.

$$\mathbf{V} \Lambda^2 \mathbf{V}^T = \mathbf{V}_{(0)}^{-\frac{1}{2}} \mathbf{S} \mathbf{V}_{(0)}^{-\frac{1}{2}} \quad (25)$$

The size of the matrix in Eq. (25) is determined by the number of outputs, so there is a total of p possible modes. To reduce the number of eigenvectors to the desired m modes, the eigenvectors associated with the m largest eigenvalues are used. The remaining eigenvectors are dropped. This initial guess is the principal component solution. The use of PCA will give the shapes that capture as much of the variations seen in the data as possible.

The initial scaling matrix, Eq. (26), is calculated from the correlation matrix and mode shapes, depending on if the mode shapes are fixed. For fixed mode shapes, the $\mathbf{P}_{(0)}$ is re-scaled to make the diagonal of one. For free mode shapes, the scaling matrix is simply calculated from the decomposition of Eq. (25).

$$\kappa_{(0)} = \begin{cases} \text{diag } \mathbf{P}_{(0)} & \text{fixed mode shapes} \\ \Lambda & \text{otherwise} \end{cases} \quad (26)$$

B. Rotations of Shapes

Noted: in general, the shapes estimated by Eq. 21a, $\hat{\Phi}$, are not mode shapes and can be rotated by an arbitrary rotation matrix, $\mathbf{T}$, to give a different set of shapes. The rotation is demonstrated and is shown in Eq. (27a). The condition for the rotation matrix is shown in Eq. (27b).

$$\bar{\Phi} \bar{\kappa} = \hat{\Phi} \kappa \mathbf{T} \quad (27a)$$

$$\text{for any } \mathbf{T} \text{ such that } \mathbf{T}^T \mathbf{T} = \mathbf{I} \quad (27b)$$

In factor analysis, these rotations are often used to try to transform the shapes into a form that is easier to interpret. The discussion of these rotations is often one of the most ambiguous and contentious parts of the analysis. The selection of a factor rotation can become somewhat of a fishing exercise to find the rotation that shows the shape the researcher is looking for. However, in engineering problems there are a few broadly accepted conditions that the mode shapes should meet. Additionally, each definition for the mode shapes has an unambiguous purpose. The challenge is mathematically defining the conditions for the rotation, rather than the selection of the type. If the shapes are estimated by the factor analysis process, they will not be typical structural mode shapes. To make these estimated shapes more representative of structural modes, a set of target mode shapes are taken from the FEM.

The current implementation of the fitting allows one to constrain the shape definitions to a set of known mode shapes. If the mode shapes are constrained, the shapes are exactly the mode shapes, and the rotation matrix becomes unnecessary. However, in the course of the present work, it became apparent that even if the mode shapes are known, it is sometimes better to fit the factor model with unknown mode shapes, and then rotate these to best match the known shapes. Allowing varying shapes helped reduce the effect of errors in the mode shapes. The mode shapes from a finite element model will never match exactly to that of an actual vehicle. Furthermore, the shapes can change slightly

with mass conditions (e.g., fuel or cargo) or with wear and fatigue. Allowing the model fitting to have varying shapes reduced the spatial correlation of the errors along the span, as discussed in section VI, “X-56A Results.” Allowing the shapes to be estimated also reduced the overfitting of a mode that was not well excited by the test point.

The rotation will change the scaling matrix. To calculate the updated scaling matrix, $\bar{\kappa}$, the factor rotation, Eq. (27a), is multiplied through by the transpose of the rotated factors, $\hat{\mathbf{Z}}^T$. Because we are only interested in the scaling, consider only the diagonal terms, Eq. (28).

$$\text{diag}(\hat{\Phi}^T \hat{\Phi} \bar{\kappa}) = \text{diag}(\hat{\Phi}^T \hat{\Phi} \kappa \mathbf{T}) \quad (28)$$

Because $\bar{\kappa}$ is defined to be diagonal, it can be separated and the expression is simplified, Eq. (29).

$$\text{diag}(\bar{\kappa}) = \text{diag}(\bar{\Phi}^T \hat{\Phi} \kappa \mathbf{T}) \oslash \text{diag}(\bar{\Phi}^T \bar{\Phi}) \quad (29)$$

1. In Vacuo Mode Orthogonal Procrustes Rotation

The first rotation is to match the calculated shapes, $\hat{\Phi}$, to a constant set of known (target) mode shapes, $\bar{\Phi}$. For the present work, the known set of modes are taken from the FEM. The solution of the optimal rotation is calculated from Procrustes analysis. The algorithm uses the SVD of the normalized modes in Eq. (30).

$$\mathbf{U}_T \mathbf{S}_T \mathbf{V}_T^T \triangleq \kappa \hat{\Phi}^T \bar{\Phi} \quad (30)$$

The SVD is a numerically efficient and robust method to get the rotation matrix in Eq. (31).

$$\mathbf{T} \triangleq \mathbf{U}_T \mathbf{V}_T^T \quad (31)$$

The rotation of Eq. (31) is the rotation that minimizes the average least square error along all the fiber gratings. The rotated shapes will not exactly match the target mode shapes, Eq. (32).

$$\bar{\Phi} \bar{\kappa} \neq \hat{\Phi} \kappa \mathbf{T} \quad (32)$$

To examine the error in the rotated modes, we consider the modal assurance criteria (MAC), Eq. (33). The diagonal of the MAC is equivalent to the coefficient of determination, R^2 , of each mode shape, or equivalently the error in the strain shape relative to its magnitude.

$$MAC \triangleq \text{diag} \left(\sqrt{\bar{\Phi}^T \bar{\Phi}} \right)^{-1} \bar{\Phi}^T \hat{\Phi} \kappa \mathbf{T} \text{diag} \left(\sqrt{\mathbf{T}^T \kappa \hat{\Phi}^T \hat{\Phi} \kappa \mathbf{T}} \right)^{-1} \quad (33)$$

The poor value of the MAC does not necessarily mean a large error in the model. In some cases, it may mean that the mode was not effectively excited in the test point. If the mode is ineffectively excited, it should be removed from the analysis and the fit recalculated, or some other method of mitigating the overfitting should be applied.

2. In Vacuo Mode Oblique Scaled Procrustes Rotation

Even if the mode shapes were constrained to the target, the Procrustes rotation will not rotate to the target shapes. The Procrustes rotation cannot correct for the scaling difference between the fit and the target. To correct for this, a scaled Procrustes rotation is defined. For the scaled Procrustes, the unscaled mode shapes are used in the SVD, Eq. (34).

$$\mathbf{U}_T \mathbf{S}_T \mathbf{V}_T^T \triangleq \hat{\Phi}^T \bar{\Phi} \quad (34)$$

The scaled Procrustes rotation matrix, Eq. (35), is scaled by the factor scaling matrix, κ ; therefore, for constant modes, this means that the rotation matrix is identity.

$$\mathbf{T} \triangleq \kappa^{-1} \mathbf{U}_T \mathbf{V}_T^T \kappa \quad (35)$$

However, the scaling means that the rotation matrix is now oblique. The oblique rotation means that the correlation matrix, $\mathbf{P}$, is not diagonal and will now capture the correlations between modes.

3. In Vacuo Mode Oblique Pattern Rotation

The pattern rotation, Eq. (36), is simply the least squares estimate between the estimated shapes and target modes.

$$\mathbf{T} \triangleq \left(\hat{\Phi}^T \text{diag} \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi} \right)^{-1} \hat{\Phi}^T \text{diag} \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \bar{\Phi} \quad (36)$$

The pattern rotation does not satisfy the condition of Eq. (27b), and is therefore not a true rotation; however, it will give the closest match to the target mode shapes.

C. Estimating Generalized Coordinates

Now that the statistical properties of the noise and generalized model coordinates have been estimated, it is possible to estimate the value of the generalized modal coordinates. Three separate methods, each with different assumptions, have been examined.

1. Ordinary Least Squares (No Factor Model)

The most obvious approach for estimating the generalized modal coordinates is to use an OLS approach, Eq. (37).

$$\hat{\boldsymbol{\eta}}^{ols} = \left[\begin{array}{c} \left(\Phi^T \Phi \right)^{-1} \Phi^T \\ (m \times n) \quad (m \times p) \quad (p \times n) \end{array} \right] \mathbf{Z} \quad (37)$$

The OLS approach needs no information from the factor model. As a result, application of OLS to estimate the generalized coordinates does not require estimates of the measurement noise properties. However, in many cases the OLS approach can perform quite poorly, particularly when the mode generalized coordinates are of significantly different magnitudes or are less observable.

2. Bartlett (Weighted Least Squares)

The first approach to estimating generalized model coordinates from the factor model is Bartlett's method or weighted least squares [12]. Bartlett's method, Eq. (38), finds the modal coordinate, $\hat{\boldsymbol{\eta}}$, which minimizes the error between the factor model output and the measured output, $\mathbf{z}$. The error is scaled by the output noise to standardize between outputs, which could have significantly different scales.

$$\hat{\boldsymbol{\eta}}_k = \arg \min_{(m \times n)} \frac{1}{2} \sum_k \left\| \mathbf{V}^{-\frac{1}{2}} (\mathbf{z}_k - \Phi \boldsymbol{\kappa} \hat{\boldsymbol{\eta}}_k) \right\|_2 \quad (38)$$

Bartlett's method is a weighted least squares problem. The estimates of the generalized model coordinates are given in Eq. (39).

$$\hat{\boldsymbol{\eta}}^{wls} = \left[\begin{array}{c} \left(\boldsymbol{\kappa} \Phi^T \mathbf{V}^{-1} \Phi \boldsymbol{\kappa} \right)^{-1} \boldsymbol{\kappa} \Phi^T \mathbf{V}^{-1} \\ (m \times n) \quad (m \times p) \quad (p \times n) \end{array} \right] \mathbf{Z} \quad (39)$$

The factor scaling, $\boldsymbol{\kappa}$, represents the variance, or root mean square (RMS), of the generalized modal coordinates. If the RMS of these modal coordinates and the measurement noise, $\mathbf{V}$, are all equal, Bartlett's method becomes identical to the OLS solution. So, one can see that the OLS solution would be a reasonable choice if the RMS of the generalized coordinates are of a similar magnitude, and the output noise is similar: such as from the FOSS. However, if the RMS of the one modal coordinate is of a smaller magnitude, such as one not effectively excited by the maneuver, the ordinary least squares can have a significantly larger error.

Bartlett's method is an effective method for estimating the generalized model coordinates. Its similarity to OLS makes it a familiar and popular choice. Bartlett's method is the minimum variance (i.e., minimum error) unbiased estimator of the generalized model coordinates [12]. However, there are biased estimators of the generalized coordinates that can have less error. Because Bartlett's method is minimizing the error in the output, the measurement noise will cause more variance in the modal coordinate estimates.

3. Regression (Thomson)

The final method being used is the regression method [12]. The regression method is based on the joint probability distribution of the outputs, $\mathbf{z}$, and the generalized model coordinates, $\boldsymbol{\eta}$. The joint probability is normal with a mean in Eq. (40a) and a covariance in Eq. (40b). The definition of the factor model, Eq. (1), means that the covariance of the outputs and the generalized model coordinates is defined by the scaled mode shapes, $\Phi\boldsymbol{\kappa}$.

$$\begin{bmatrix} \boldsymbol{\mu}^* \\ (m+p) \times 1 \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{z}} \\ \mathbf{0} \end{bmatrix} \quad (40a)$$

$$\begin{bmatrix} \boldsymbol{\Sigma}^* \\ (m+p) \times (m+p) \end{bmatrix} = \begin{bmatrix} \Phi\boldsymbol{\kappa}\mathbf{P}\boldsymbol{\kappa}\Phi^T + \mathbf{V} & \Phi\boldsymbol{\kappa} \\ \boldsymbol{\kappa}\Phi^T & \mathbf{P} \end{bmatrix} \quad (40b)$$

The maximum likelihood estimate of the modal coordinate is therefore the expected value of the modal coordinate given (conditional on) the output, $E(\boldsymbol{\eta}|\mathbf{z})$. The mean and covariance of the conditional distributions are shown in Eqs. (41) and (42).

$$E(\boldsymbol{\eta}|\mathbf{z}) = \underset{(m \times p)}{\boldsymbol{\kappa}\Phi^T} \left(\underset{(p \times p)}{\Phi\boldsymbol{\kappa}\mathbf{P}\boldsymbol{\kappa}\Phi^T + \mathbf{V}} \right)^{-1} \underset{(p \times n)}{(\mathbf{Z} - \bar{\mathbf{z}})} \quad (41)$$

$$E(\boldsymbol{\eta}^T \boldsymbol{\eta}|\mathbf{z}) = \underset{(m \times m)}{\mathbf{P}} - \underset{(m \times p)}{\boldsymbol{\kappa}\Phi^T} \left(\underset{(p \times p)}{\Phi\boldsymbol{\kappa}\mathbf{P}\boldsymbol{\kappa}\Phi^T + \mathbf{V}} \right)^{-1} \underset{(p \times m)}{\Phi\boldsymbol{\kappa}} \quad (42)$$

Equation (41) is rearranged using the Sherman-Morrison-Woodbury (matrix inversion) Lemma [16] to give the regression estimate of the modal coordinate, Eq. (43). This implementation of Eq. (43) is helpful because it is inverting an $m \times m$ matrix which is much smaller and therefore faster and more well numerically conditioned.

$$\underset{(m \times n)}{\hat{\boldsymbol{\eta}}^{reg}} = \left[\left(\underset{(m \times p)}{\mathbf{I} + \mathbf{P}\boldsymbol{\kappa}\Phi^T \mathbf{V}^{-1} \Phi\boldsymbol{\kappa}} \right)^{-1} \underset{(p \times n)}{\boldsymbol{\kappa}\Phi^T \mathbf{V}^{-1}} \right] \underset{(p \times n)}{\mathbf{Z}} \quad (43)$$

By rearranging Eq. (43), Bartlett's method can be written as a function of the regression solution, Eq. (44).

$$\left[\left(\underset{(m \times m)}{\boldsymbol{\kappa}\Phi^T \mathbf{V}^{-1} \Phi\boldsymbol{\kappa}} \right)^{-1} + \underset{(m \times m)}{\mathbf{P}} \right] \underset{(m \times n)}{\hat{\boldsymbol{\eta}}^{reg}} = \underset{(m \times n)}{\hat{\boldsymbol{\eta}}^{wls}} \quad (44)$$

If the output noise is very small, $\mathbf{V} \rightarrow 0$, and the signals are uncorrelated, $\mathbf{P} = \mathbf{I}$, the regression solution will match the Bartlett's method solution. Because the regression solution is the maximum likelihood estimate of the generalized model coordinates, it is expected that, in general, it will be a lower error solution than Bartlett's method. However, the regression solution is a biased estimate of the generalized model coordinates. The bias means that the regression estimator will typically give factor scores that have a variance smaller than the true generalized coordinates. Although Bartlett's method was discussed first here, the Bartlett's method was actually developed later as an unbiased estimator. However, for the present purposes, the lower error in the estimates is preferable to the unbiased (or less biased) estimates of Bartlett's method.

It should be noted that the regression estimation is analogous to a Kalman filter, where the frequencies of the poles are infinite. The similarity is apparent if one considers that they are both maximum likelihood estimators of the state. The primary difference is that the Kalman filter also requires the differential equation of the generalized coordinates.

4. Other Methods of Estimation

There are numerous other methods for estimating the factor scores; however, the presented methods are the most common approaches. As shown in the results, both Bartlett's method and the regression methods give very similar results. It is not expected that the other methods would give significantly better results, and any better result is likely to be coincidental. Furthermore, the orthogonal factor model is much more common than the oblique. As a result, most methods of estimating the factor scores are specialized to the orthogonal factor model. We determined that it was not worth attempting to extend these methods to the oblique case without a clear benefit.

VI. X-56A Results

The X-56A MUTT, Fig. 3, was developed specifically to study active control of the structure of an aircraft and suppression of body freedom flutter [17, 18]. The wings of the aircraft can be swapped to fly an aircraft with a relatively stiff structure or a flexible aircraft with unstable flutter modes. The primary flutter mode of the X-56A MUTT is a body freedom flutter mode, the mechanism of which is the short-period being driven unstable by the structural dynamics. Unlike most production aircraft today, the flutter mode is within the flight envelope. The X-56A MUTT also offers an unusual case of operating at flight speeds faster than open-loop flutter. The use of a flutter suppressing control law allowed flight operations to airspeeds past where the open-loop flutter mode would go unstable. The X-56A MUTT has a wingspan of 28 ft and a designed flight speed of 50 to 150 knots. However, the maximum level speed was approximately 130 knots, and the flex-wing tests went to a maximum airspeed of 120 knots. The flutter speed was approximately 105 to 115 knots depending on the fuel state. The X-56A MUTT has a relatively low subsonic flight envelope. The absence of transonic effects does not limit the applicability of the technology application because the present work is focused on the structural aspects. It is important to demonstrate that any estimation method is robust to such modeling errors that will be present on a flight vehicle. The X-56A MUTT was equipped with the FOSS for 22 of the 31 flex-wing flights. The FOSS on the X-56A MUTT used optical frequency domain reflectometry [19]. As a result, there is a large number of strain measurements, but a relatively low sampling rate capability.

A. Description of Experiment

The X-56A flight test was selected as the primary example for study. A detailed finite element model of the X-56A MUTT is available in NASTRAN™ (MSC Software Corporation) (Newport Beach, California) [20]. There were also aeroelastic models [21] available that were able to accurately model the aeroelastic modes of the X-56A MUTT. A total of four fibers were used on X-56A MUTT. The layout of the fiber-optic sensing on the upper and lower surfaces of the left wing of the X-56A MUTT is shown in Fig. 4. The fiber was laid out symmetrically on the right wing.




Fig. 4 X-56A left wing fiber-optic layout in blue.

The zig-zag patterns were intended to better capture torsional strain, but because of complexities in locating the gratings, and a poor signal to noise, these portions of the fiber have not been used in the current analysis. The grating locations can be slightly different between the left and right wings; however, this effect was ignored, and the fiber layout was assumed to be symmetric from left to right. The upper and lower surfaces are very similar in layout, but there are slight differences in where the gratings end up being located because of a need to accommodate the actuator access panels on the lower surface. The arrangement of the four fibers provided a total of 936 strain measurements at a sampling rate of 40.711 Hz. It may be possible to identify dynamics up to ~ 20 Hz; however, in practice [22] the dynamics will only be accurate to ~ 8 Hz. As a result, it is difficult to use traditional modal analysis tools to get any mode shapes other than the first symmetric and antisymmetric bending for the X-56A aircraft.

Two test points, Table 2, from flights 16 and 26 of the X-56A vehicle were selected for the case study.

Table 2 X-56A test case flight conditions.

Case	Airspeed, kn	Fuel weight, lb	Flight	FTA
1	80	62	16	506
2	110	56	26	5507

Both of these flight conditions were using the classical control architecture. The first condition was a low-speed case where the structural modes were well damped and before there was significant aerodynamic coupling. The second condition was a much higher airspeed and was past the open-loop flutter speed. Both conditions were using a similar multisine input. The left and right surfaces were excited in symmetric pairs to excite the longitudinal dynamics; however, some lateral dynamics were excited as a consequence of asymmetries in the vehicle. They were both multisine maneuvers designed to excite frequencies from ~ 0.2 to ~ 16.0 Hz. The frequency range covers the three primary modes associated with the body freedom flutter, short-period, first symmetric bending, and first symmetric torsion. The only difference between the maneuvers was the magnitude of the inputs. The higher dynamic pressure required a smaller input to avoid exceeding operational limits.

A description and the acronyms of the flex modes in the FEM are shown in Table 3. As a result of a preload in the nose gear springs, the nose gear did not deflect in flight. Also, the nose gear stroke produced no strain in the wing to be observed by the FOSS; therefore, the nose gear mode was not used in any of the analyses. The *in-vacuo* frequencies are included for reference; the flight-frequencies do not match these, but we can see that the selected multisine maneuvers are expected to excite up to mode eight.

Table 3 X-56A FEM mode shapes.

Mode number	Acronym	Frequency at 60 lb fuel, Hz	Description
1	SW1B	3.06	Symmetric wing 1st bending
2	AW1B	5.20	Antisymmetric wing 1st bending
3	NG	8.71	Nose gear stroke
4	SW1T	11.11	Symmetric wing 1st torsion
5	SFA	12.17	Symmetric wing fore-aft
6	AW1T	12.45	Antisymmetric wing 1st torsion
7	SMLG	13.12	Symmetric main landing gear
8	AMLG	13.96	Antisymmetric main landing gear
9	SWL	15.83	Symmetric winglet bending
10	Boom Lat	17.84	Air data boom lateral
11	AW2B	18.01	Antisymmetric wing 2nd bending
12	Boom Vert	18.23	Air data boom vertical

B. Maximum Likelihood Factor Analysis of Simulated Data

The maximum likelihood factor analysis method was first applied to simulated data from the X-56A linear models. Because simulated data are used, the exact solution of the modes and the generalized coordinates are available for comparison. An identically and independently distributed (IID) Gaussian random noise was applied to all of the gratings. The noise magnitude was selected to be similar to the flight data. The simple Gaussian noise would not capture the complexity of the noise behavior seen in the X-56A ground proof testing; however, the simulated data was useful for assessing the different proposed methodologies and identifying potential problems.

A fit with a total of six modes was performed. The higher order fit included all modes that were expected in the frequency bandwidth of the maneuver. The number of modes in the higher order fit may have been excessive, but it was useful to see how the method degraded as extra modes were added. In an effort to create a better fit, the MAC/R^2 error

metric of the modes were then used to remove modes that were not well captured by the test data. The AW1B mode was also retained because it was the only antisymmetric mode that remained. The antisymmetric mode was needed when looking at the flight data because the FEM may not have captured all of the asymmetries present in the vehicle. The reduction of modes resulted in a total of six modes remaining. The fit was performed with a variety of different settings to demonstrate the potential benefits of each of the different methods that have been developed. These settings are shown in Table 4.

Table 4 Simulation fitting cases.

Case Name	Mode Shapes	Rotation	State Estimator	Sensor Noise
Exact	N/A	N/A	N/A	N/A
OLS	Constant	N/A	OLS	Free
FA Regression	Varying	Procrustes	Regression	Free
FA Bartlett	Varying	Procrustes	Bartlett	Free
FA Pattern	Varying	Pattern	Regression	Free
FA Constant Modes	Constant	N/A	Regression	Free
FA IID Noise	Constant	N/A	Regression	Uniform

1. Modal Coordinate Root Mean Square

The first step in the estimation process is to fit the factor model of Eq. (1). This step provides an independent estimation of the variance of the generalized model coordinates, without actually estimating their value at each time step. To assess the effectiveness of the factor model fitting, these variances for the X-56A low-speed case are shown in Fig. 5. The regression and Bartlett methods are identical in this case. The response is dominated by the SW1B mode, so all the methods are good at estimating that mode. Other than the pattern rotation case, the methods do provide a good estimate for SW1T. In general, the pattern rotation does a poor job estimating the variances. The cases with constant modes work well when the number of modes is six; however, the constant mode shapes performed very poorly with higher order fits. The IID noise case shows poorer performance, and it makes little difference when reduced to six modes. The OLS does give the best estimates of variance. This is partially because all simulated sensors do have identical noise characteristics. However, when increasing to the number of modes, the OLS rapidly degrades and significantly overestimated the mode state variance for modes of smaller magnitude. The regression method shows a much better estimate of the mode state variances for the smaller modes such as SFA and AW1T. Overall, this suggests that for a well formulated mode estimation, the OLS will often perform very well. However, if small magnitude or higher order modes are needed, the present factor analysis methods are more effective.




Fig. 5 Estimates of variance of simulated generalized modal coordinates, low speed.

The X-56A high-speed case, Fig. 6, shows results very similar to the low speed. The SW1B mode works well for all methods, and the SMLG works poorly for all methods. The OLS works well but tends to overfit if additional modes are added. The pattern rotation gives poor estimates of the variance.




Fig. 6 Estimates of variance of simulated generalized modal coordinates, high speed.

2. Output Noise Estimates

The other component of the factor model of Eq. (1) is the value of the sensor noise. In Fig. 7 a box plot is used to show the variations in the estimates of the noise compared to the analytical value used in the simulation. In the box plot, the box shows the range of 50% of the data. The black whiskers show the range that one would typically expect any remaining data to lie within. Any measured points outside this range, *outliers*, are shown by red marks. In this case, the fit for the regression, Bartlett, and pattern rotation cases are identical. The rotation has no effect on the estimates of the sensor noise. These cases with the mode shapes as varying parameters were able to capture the sensor noise extremely well. The IID noise case has no dispersion by definition. It is overestimating the noise in Fig. 7. The constant modes without the IID noise are typically overestimating the noise and have an excessively large dispersion that is literally putting the median off the chart. In Fig. 7, the OLS, which did a better job of estimating the mode state variance, is significantly overestimating the noise. This overestimation is caused by spillover of the higher order modes not in the estimation. The difficulty of OLS to correctly estimate higher order modes, means that it is unable to give an accurate estimate of the sensor noise.




Fig. 7 Estimates of sensor noise variance for simulated data, low speed.

The box plot of the output noise estimates for the high-speed case are shown in Fig. 8. The regression, Bartlett, and pattern rotation cases capture the sensor noise very well. The constant mode case is again giving unusable estimates of the sensor noise. The constant modes with the IID noise constraint show the same decreasing sensor noise with the number of modes. The noise estimates from OLS are still more accurate for the higher order fit in Fig. 8. However, the higher dynamic pressure means that there is more aerodynamic coupling between modes and therefore more mode spillover. The additional spillover is further degrading the mode estimates, making the OLS less useful in estimating the noise.




Fig. 8 Estimates of sensor noise variance for simulated data, high speed.

The problems with the noise estimates from the OLS and constant mode are further expanded by looking at the noise estimates along the length of the fibers in Fig. 9. The constant modes show excessive overestimation of the noise. The constant modes also show correlation along the fiber, i.e., a clear pattern, in the errors. Forcing a IID noise does help address the correlations in the noise, and it does provide reasonable estimates of the noise; however, still not as accurately as the varying modes. The OLS noise estimates, do not show as many problems as the constant modes, but they do show a slight pattern. This pattern is again caused by spillover of higher order modes.




Fig. 9 Noise estimates along the fiber for simulated data, low speed.

Looking at the noise estimates along the fiber for the high-speed case, Fig. 10, shows exactly the same problems for the constant modes and OLS methods.




Fig. 10 Noise estimates along the fiber for simulated data, high speed.

If one wishes to get an accurate estimate of the sensor noise, the factor analysis method with varying modes is providing clearly superior results. The OLS can converge to the correct result if sufficient modes are included, but the large number of modes required means that state estimates would therefore be useless. The constant mode case has a severe problem with erroneous correlation in the noise estimates along the fiber. Forcing the noise values to be uniform does improve the estimates but is still not as accurate as the varying mode case. Furthermore, the sensor noise may not be uniform for real flight data, especially if different sensors are used.

3. Goodness of Fit of Mode Shapes

The final component to consider in the factor model is the effectiveness of the different rotation options. The coefficient of determination of the mode shapes of the simulation data from the different rotations are compared in Fig. 11 for the low-speed case. The constant shapes use the exact target, so they match perfectly with a value of one. As a rotation matrix, the Procrustes is more constrained and has a lower coefficient of determination than the pattern, but for most of the mode shapes, the Procrustes and pattern rotations have similar goodness of fit. The AW1B mode was retained because it had the best fit of the antisymmetric modes. It was found that the small vehicle asymmetries created problems

if there were no antisymmetric modes, especially for the flight data. The Procrustes shows a poor fit for the SMLG mode, but because the pattern fit was good it was retained for later analysis. Based on the later results, it appears that it may have been better to have removed the SMLG mode as well. Although the constant and pattern rotations did present with the best mode shape fits, these methods returned the poorest results for the other aspects of the model. Further development of these rotations is likely needed.




Fig. 11 Similarity of mode shapes from simulated data, low speed.

The coefficient of determination of the mode shapes for the high-speed case are shown in Fig. 12. Overall, the three different rotation types show very similar effectiveness for the high-speed case. The AW1T and AMLG modes do show a slightly poorer fit for the high-speed case.




Fig. 12 Similarity of mode shapes from simulated data, high speed.

The Procrustes and pattern rotations do show similar overall goodness of fit. The pattern rotation does show a slightly better fit to the SMLG mode. The constant modes obviously match exactly by definition; however, given the problems seen with the sensor noise estimates and the mode state variance, the Procrustes method is expected to give significantly better overall performance in estimating the generalized model coordinates.

4. Modal Coordinate Estimates

After fitting the factor model of Eq. (1), it possible to construct the estimates of the generalized coordinates. For the simulated data, we can compare to the exact generalized model coordinates from the simulation. The ability to compare

to the exact data gives a good indication of the ability and limitations of the different estimators. An OLS solution was also used to compare to a more traditional approach. The confidence bounds of the OLS estimates are also shown as a shaded region. A small arbitrary time slice was selected to more clearly show the fit of the methods. The modal coordinates for the X-56A high- and low-speed cases are very similar; therefore, only the high-speed cases will be examined in detail. For the estimate of SW1B, Fig. 13, all of the methods are giving a very accurate estimate. The constant mode solution is slightly underestimating some of the peaks but is still a good estimate of SW1B.




Fig. 13 Time history of estimate of simulated SW1B, high speed, six total modes.

For the AW1B mode, Fig. 14, the constant mode solution is showing a tendency to be overfitting. All of the other methods give a good estimate of the exact solution.




Fig. 14 Time history of estimate of simulated AW1B, high speed, six total modes.

In the estimates of SW1T, Fig. 15, one can see that the mode state estimates are consistent between all of the methods and are very close to the exact solution. Removing the extra modes has significantly improved the accuracy of the OLS, Bartlett, and IID noise estimators. The other methods were already giving a good estimate.




Fig. 15 Time history of estimate of simulated SW1T, high speed, six total modes.

In the estimates of SFA, Fig. 16, the OLS, Bartlett, and IID noise estimators are closer to the exact. These solutions are not giving results identical to the pattern rotation solution; however, the regression and constant mode cases are still giving better estimates.




Fig. 16 Time history of estimate of simulated SFA, high speed, six total modes.

In the estimate of SMLG, Fig. 17, the OLS, Bartlett, and IID noise estimators are closer to the exact solution. Because all the methods are close to the exact, the errors in the regression and constant mode cases are clearer. The regression and constant modes show an overestimation of the modal variations.




Fig. 17 Time history of estimate of simulated SMLG, high speed, six total modes.

Figure 18 shows all of the methods are giving a good estimate of the SWL mode in.




Fig. 18 Time history of estimate of simulated SWL, high speed, six total modes.

By removing antisymmetric modes that are not well excited by the maneuver, the accuracy of the OLS, Bartlett, and IID noise estimators can be improved significantly. The partitioning of modes into symmetric and antisymmetric is a common and effective tool in postflight analysis; however, in a real-time or flight control application, it may not be possible to know *a priori* which modes to remove. As was seen with the AW1B, the vehicle may be more asymmetric than predicted, or the maneuvers may not be as symmetric as designed. Furthermore, the ability to accurately estimate the higher frequency modes give a more accurate estimate of the actual vehicle shape.

C. Maximum Likelihood Factor Analysis of Flight Data

To apply all the methods to the flight data, there were too many results to be easily reviewed. So, only the best performing method was used, and the simulated and flight-test results have been compared. With the simulated data, the fitting of the factor model worked best with the varying modes and the scaled Procrustes rotation. For estimating the mode generalized coordinates, the regression method was selected because generally this method showed less sensitivity to the noise. The same six modes used with the simulated data have been used with the flight test data.

1. Modal Coordinate Root Mean Square

Figure 19 shows the estimates of the variance from the simulated and flight data at low speed. As seen previously, the simulation is very accurate at capturing the SW1B, SW1T, and SFA modes. The flight data shows very similar estimates of variance for these three modes. The flight data variances are generally slightly smaller. Although the simulation data does not estimate the SMLG and SWL modes well, the simulated fits and flight data are close. This result suggests that the issue is primarily the signal to noise and is well replicated in the simulated data. The AW1B mode does show significantly different behavior between the flight and simulated data. With the fit in Fig. 19, the AW1B mode is the only antisymmetric mode; any antisymmetric dynamics are being fit to that mode. The lack of additional antisymmetric modes means that the analysis is overfitting the AW1B mode. However, because the maneuver is symmetric and was not designed to excite the antisymmetric modes well, it is impractical to add higher order antisymmetric modes to get a better fit. The flight data is also likely to be larger magnitude because the composite structure of the vehicle is likely to have more asymmetry than was captured in the FEM. The asymmetry of the wings resulting from the composite manufacturing was observed previously when the wings underwent modal testing individually on a strongback, but the difference was small and considered impractical for any model adjustment to capture that effect.




Fig. 19 Estimates of variance of flight generalized modal coordinates, low speed.

The estimator shows similar performance with the flight data for the high-speed case in Fig. 20. The high-speed case does show more excitation of the antisymmetric modes; however, the larger magnitude is seen in the exact results as well and is likely caused by increased aerodynamic coupling.





Fig. 20 Estimates of variance of flight generalized modal coordinates, high speed.

2. Output Noise Estimates

The statistics of the noise estimates from the flight data are shown in Fig. 21. The box plot of the noise in Fig. 21a shows that the median noise along the fiber was slightly lower than the value used in the simulation; however, it is within the interquartile range, so it is likely to be a statistical anomaly. The flight data does show much more variation in the estimates than the simulated data. The flight data also shows a lot more outliers (shown as red Xs). The dependence of the noise characteristics on the loading can produce these outliers; however, looking at the variations of the grating noise along the fiber in Fig. 21b, there is not a clear pattern in these outliers. The lack of a pattern indicates that the method is effectively capturing the known outliers in the data and has no spillover from higher modes.



a) Dispersion of noise estimates.

b) Along all fibers.

Fig. 21 Noise estimates along the fiber for flight data, low speed.

The sensor noise estimates for the high-speed case are shown in Fig. 22. The median noise level from the high-speed case was used for the simulation, so the box plot, Fig. 22a, shows an exact match between the simulation and flight. Because the loading is different, the gratings which are showing the larger outlier values of noise in Fig. 22b are different than the low speed; however, there is still not a clear pattern in these residuals. Otherwise, the low- and high-speed cases show very similar noise estimation characteristics.




Fig. 22 Noise estimates along the fiber for flight data, high speed.

3. Goodness of Fit of Mode Shapes

The coefficient of determination for the mode shapes are compared between the simulation and the flight data in Fig. 23. The AW1B shows a much higher R^2 in the flight data. This implies that the mode was much more prevalent in the flight data than the simulated data. The SFA mode shows a much lower R^2 for the flight data versus the simulated results. The accuracy of the variance estimate, Fig. 19, suggests that the mode was well excited, but that the mode shape was likely different.




Fig. 23 Similarity of mode shapes from flight data, low speed.

The goodness of fit of the mode shapes for the high-speed case is shown in Fig. 24. The high-speed results, Fig. 24, are almost identical to the low-speed results, Fig. 23.




Fig. 24 Similarity of mode shapes from flight data, high speed.

4. Modal Coordinate Estimates

The X-56A flight data has additional turbulence noise that makes it difficult to compare the simulation and flight time histories. When working with multisines, time synchronization of the two different data types is difficult. Instead of time histories, the power spectral densities have been used to compare the simulation and flight data. The power spectral estimates are very jagged because of the multisine. The multisine excite each symmetric pair of surfaces at individual frequencies. The differences in control effectiveness mean that the magnitude between frequencies can be quite different. Because this variation is a physical effect, and not simply noise, smoothing of these power spectra to remove the jaggedness is not appropriate.

CASE 1, LOW SPEED The variations with frequency do appear to be smaller for the flight data; however, this may be an artifact of the uneven sampling rate of the FOSS system. The simulation data needed to run at 40 Hz to correspond to an even multiple of the flight-test inputs at 200 Hz. Because the multisine were designed for the 200-Hz sampling rate, the FFT frequencies line up well with the simulation data. However, the FFT of the flight data has frequencies that do not line up with the multisine frequencies, and this inconsistency will have a smoothing effect to the frequency response for the flight data.

As seen with the simulated data, the first mode in Fig. 25 matches very well between the exact, the simulation fit, and the flight-test data. The flight data does not have the step at the edge of the multisine because it has the additional excitation from turbulence.




Fig. 25 Power spectrum estimate of flight SW1B, low speed, six total modes.

In Fig. 26, the AW1B shows a significantly higher power from the flight data, as previously shown in Fig. 19. However, the flight data does clearly show the peak from the AW1B mode at ~ 5 Hz. The simulation fit was not able to capture the peak from the AW1B mode. This result suggests that the asymmetry in the vehicle is much larger than captured in the model.




Fig. 26 Power spectrum estimate of flight AW1B, low speed, six total modes.

Figure 27 shows the power of the SW1T mode showing good agreement between the simulated results and the flight-test results. The flight data does show a slightly lower magnitude at frequencies past the first wing bending. Given the similarity in shapes and frequencies, the difference may be a matter of the control effectiveness in the simulation.




Fig. 27 Power spectrum estimate of flight SW1T, low speed, six total modes.

Figure 28 shows good agreement between the simulation and the flight data show for the SFA mode. The flight data does show a higher peak and a slightly higher frequency for the SFA mode. The fore-aft motion of the wing is a challenge for the doublet lattice aerodynamics to capture.




Fig. 28 Power spectrum estimate of flight SFA, low speed, six total modes.

The flight data and the simulation fit show very similar responses in Fig. 29. The simulation data does show significant errors relative to the exact. However, the fact that the simulated and flight data are similar suggests that the simulated data is effectively reproducing the problems from noise seen in the flight data. It would probably be beneficial to remove the SMLG mode from a future analysis, but in this case, the mode is helpful for assessing the capabilities of the new method.




Fig. 29 Power spectrum estimate of flight SMLG, low speed, six total modes.

Similarities in the overall gain of the SWL mode for the simulated and flight data are shown in Fig. 30. However, the flight data does show a peak in the SWL mode at the AW1B mode that is not replicated by the simulated data. The flight data is showing a large peak at the SW1T mode suggesting that the simulation does not capture the coupling between the SWL and the SW1T. The edge of the multisine makes it difficult to interpret the SWL mode peak, but the flight data clearly has a much lower magnitude.




Fig. 30 Power spectrum estimate of flight SWL, low speed, six total modes.

CASE 2, HIGH SPEED As with all the other test cases, the SW1B mode shows a good match for the high-speed case, as shown in Fig. 31. There does appear to be a more apparent bump at the SW1T frequency for the high airspeed although the bump is likely a result of more aerodynamic coupling.




Fig. 31 Power spectrum estimate of flight SW1B, high speed, six total modes.

The results for the AW1B in the high-speed case, Fig. 32, show a similar larger response in the flight-test results; however, the flight-test data does appear to be slightly closer to the simulation for the higher airspeed.




Fig. 32 Power spectrum estimate of flight AW1B, high speed, six total modes.

The SW1T mode from the high-speed case in Fig. 33 does show a larger difference between the simulation and flight-test data than was seen from the low-speed case, Fig. 27. Additionally, there is a small shift in the frequency of the SW1T mode between the model and the flight for the higher airspeed.




Fig. 33 Power spectrum estimate of flight SW1T, high speed, six total modes.

The SFA mode for the high-speed case, Fig. 34, does not show a peak in the flight data that corresponds to the SW1T mode. This indicates that there is significantly less coupling between the SFA and SW1T modes at the higher airspeed, and the model does not predict the effect.




Fig. 34 Power spectrum estimate of flight SFA, high speed, six total modes.

The estimation still does not work well for the SMLG mode at the high-speed condition, shown in Fig. 35. The simulation and flight data fits are still very similar, which indicates that the simulation is still replicating the effectiveness of the fit at the higher airspeed.




Fig. 35 Power spectrum estimate of flight SMLG, high speed, six total modes.

The PSD of the SWL mode for the high-speed case are shown in Fig. 36. The SWL mode shows the largest discrepancies between the simulation and flight data. The problems with the simulation are not surprising given the difficulty in capturing the aerodynamic interference between the wing and winglet that would be critical for this mode. The simulation predicts a much larger, approximately three times, excitation of the SWL by the control system for this maneuver. The simulation also shows a much lower magnitude at the SWIT frequency but may be caused by the lower excitation of the winglets.




Fig. 36 Power spectrum estimate of flight SWL, high speed, six total modes.

The modal filtering applied to the flight data shows very similar performance to the simulated data. We can therefore infer that the flight data shows similar performance as seen with the simulation data. Furthermore, the application to the flight data was able to show a number of deficiencies in the model that were not able to be identified using more classical modal approaches while looking at the frequency response functions without estimating the generalized model coordinates.

VII. Conclusions

A new maximum likelihood method has been developed for estimating aircraft generalized model coordinates from a large number of distributed strain sensors. The method was applied to simulated and flight-test data from the X-56A flex-wing configuration.

The simulated data was used to demonstrate the accuracy of the solution. The method also worked very well with the flight-test data and was able to demonstrate deficiencies in the model that were indiscernible using other methods. The factor model was used to estimate the statistics of the data and was able to more accurately predict the noise than traditional engineering approaches such as ordinary least squares. Furthermore, the new maximum likelihood factor analysis method was shown to be capable of handling models with more generalized coordinates than least squares, which was prone to overfitting. The factor model estimation was done with constant and varying mode shapes; however, the constant mode shapes had many problems in the solution. Additional work is needed to determine if these problems are a result of local minima that can be addressed or an inherent limitation of the method. Instead, a method of estimating with varying mode shapes, and then rotating to the target mode shapes was found to give much better results.

Numerous different methods from the literature [12, 23] were considered for the state estimation step. These methods were compared against an ordinary least squares solution – a more classical approach. In general, the regression method is giving the best estimates. However, as Thomson, the inventor of the regression method noted, it is not the best method for all applications.

“Bartlett’s estimates and the regression estimates attain different ends, and it is agreed that each method is correct in the right place.” [23]

Work in other fields [24] has shown that the regression estimators can introduce bias when used in the output in system identification. Further understanding of these potential biases is needed before these methods should be applied to system identification or other specific applications. The present method of state estimation exploits the large number of strain sensors and does not require a differential equation or the system inputs to estimate the generalized coordinates, as a Kalman filter would. Furthermore, it is dependent only on the current time and requires no time to converge.

Acknowledgments

This work has been funded by the NASA Space Technology Mission Directorate through the center innovation fund. The X-56A flight-test data was available from flight tests supported by the Air Force Research Lab, the NASA Advanced Air Transport Technology Program, and the NASA Flight Demonstrations and Capabilities Program.

The authors thank all the people who contributed to the X-56A flight test. As well, the authors thank the various interns that have supported the X-56A research, specifically, Loren Newton, Noah Edwards, and Jack Button whose work helped to identify the problems that motivated the current work.

References

- [1] Love, M. H., Zink, P. S., Wieselmann, P. A., and Youngren, H., “Body Freedom Flutter of High Aspect Ratio Flying Wings,” AIAA Paper 2005-1947, April 2005.
doi: 10.2514/6.2005-1947
- [2] Noll, T. E., Brown, J. M., Perez-Davis, M. E., Ishmael, S. D., Tiffany, G. C., and Gaier, M., “Investigation of the Helios Prototype Aircraft Mishap, Vol. 1 - Mishap Report,” 2004.
- [3] Meirovitch, L., *Fundamentals of Vibrations*, McGraw-Hill, New York, 2001.
- [4] Suh, P. M., Chin, A. W., and Mavris, D. N., “Virtual Deformation Control of the X-56A Model with Simulated Fiber Optic Sensors,” AIAA Paper 2013-4844, Aug. 2013.
doi: 10.2514/6.2013-4844
- [5] Grauer, J. A., “Method for Real-Time State Estimation of Structural Modes for an Aeroelastic Wind Tunnel Model,” AIAA 2021-1643, Jan. 2021.
doi: 10.2514/6.2021-1643
- [6] Martins, B. L., and Kosmatka, J. B., “Health Monitoring of Aerospace Structures via Dynamic Strain Measurements: An Experimental Demonstration,” AIAA Paper 2020-0701, Jan. 2020.
doi: 10.2514/6.2020-0701
- [7] Shames, I. H., and Dym, C. L., *Energy and Finite Element Methods in Structural Mechanics*, Taylor & Francis Books, Inc., Great Britain, 2003.
- [8] Johnson, R. A., and Wichern, D. W., *Applied Multivariate Statistical Analysis*, Pearson Prentice Hall, 2007.
- [9] Bisplinghoff, R. L., Ashley, H., and Halfman, R. L., *Aeroelasticity*, Dover Publications, Inc., New York, 1996.
- [10] Cooley, J. W., and Tukey, J. W., “An Algorithm for the Machine Calculation of Complex Fourier Series,” *Mathematics of Computation (MCOM)*, Vol. 19, No. 90, 1965, pp. 297–301.
doi: 10.1090/S0025-5718-1965-0178586-1

- [11] Chatterjee, A., "An introduction to the proper orthogonal decomposition," *Current Science*, Vol. 78, No. 7, April 2000, pp. 808–817.
- [12] Lawley, D. N., and Maxwell, A. E., *Factor Analysis as a Statistical Method*, Butterworth & Co. (Publishers) Ltd., 1963.
- [13] Maine, R. E., and Iliff, K. W., "Formulation and implementation of a practical algorithm for parameter estimation with process and measurement noise," AIAA Paper 80-1603, Aug. 1980.
doi: 10.2514/6.1980-1603
- [14] Engle, R., and Watson, M., "A One-Factor Multivariate Time Series Model of Metropolitan Wage Rates," *Journal of the American Statistical Association*, Vol. 76, No. 376, 1981, pp. 774–781.
- [15] Jöreskog, K. G., "A General Approach to Confirmatory Maximum Likelihood Factor Analysis," *Psychometrika*, Vol. 34, No. 2, June 1969, pp. 183–202.
- [16] Rugh, W. J., *Linear System Theory*, Prentice-Hall, Inc., NJ, 1996.
- [17] Beranek, J., Nicolai, L., Buonanno, M., Burnett, E., Atkinson, C., Holm-Hansen, B., and Flick, P., "Conceptual Design of a Multi-Utility Aeroelastic Demonstrator," AIAA Paper 2010-9350, Sept. 2010.
doi: 10.2514/6.2010-9350
- [18] Miller, C. J., Schaefer, J., Boucher, M., Ouellette, J., and Howe, S., "X-56 Flight-test Approach for Envelope Expansion Past Open-loop Flutter Instability," NATO Science and Technology Organization (STO), STO-MP-AVT-211, 2022, pp. 18-1–18-27.
- [19] Richards, W. L., Parker Jr, A. R., Ko, W. L., Piazza, A., and Chan, P., "Application of Fiber Optic Instrumentation (Validation des systemes d'instrumentation a fibres optiques)," AGARD-RTO-AG-160, Vol. 22 / SCI-228, 2012.
- [20] Chin, A. W., Truong, S. S., and Spivey, N. D., "X-56A Structural Dynamics Ground Testing Overview and Lessons Learned," AIAA Paper 2020-1015, Jan. 2020.
doi: 10.2514/6.2020-1015
- [21] Ouellette, J. A., and Valdez, F. D., "Generation and Calibration of Linear Models of Aircraft with Highly Coupled Aeroelastic and Flight Dynamics," AIAA Paper 2020-1016, 2020.
doi: 10.2514/10.2514/6.2020-1016
- [22] Morelli, E. A., and Klein, V., *Aircraft System Identification: Theory and Practice, Second Edition*, Sunflyte Enterprises, Williamsburg, VA, 2016.
- [23] Thomson, G. H., "Methods of Estimating Mental Factors," *Nature*, Vol. 141, 1938, pp. 246.
- [24] Devlieger, I., Mayer, A., and Rosseel, Y., "Hypothesis Testing Using Factor Score Regression: A Comparison of Four Methods," *Educational and Psychological Measurement*, Vol. 76, No. 5, 2016, pp. 741–770.