

Gradient-Enhanced Reliability Analysis of Transonic Aeroelastic Flutter

Bret K. Stanford*

NASA Langley Research Center, Hampton, VA, 23681

Annie Sauer†

Virginia Polytechnic Institute and State University, Blacksburg, VA, 24061

Kevin E. Jacobson‡, James E. Warner§

NASA Langley Research Center, Hampton, VA, 23681

Reliability computation for transonic aeroelastic systems is a challenging task, given the complex physical mechanisms involved, and broader challenges associated with computing small probabilities of failure/flutter. This paper demonstrates the use of four different reliability estimation tools, with a particular focus on tools that can utilize gradient information: the derivative of flutter boundaries with respect to random parametric inputs. Results are shown for two test cases: the AGARD 445.6 case and the Common Research Model, and the strength and weaknesses of each UQ tool are compared in terms of cost and accuracy. Two of the UQ tools used here require gradients as part of a search process, but the other two are sample-based, and are able to leverage the existence of gradients for improved performance.

I. Introduction

The design and certification of flight structures requires the demonstration of a flutter margin, where the aeroelastic flutter boundary must remain outside the flight envelope by a 15% margin, as measured by equivalent air speed. This margin has been in use for decades and is largely empirical.¹ Furthermore, this flutter constraint is typically most aggressive in the transonic flow regime, where flutter dynamic pressure q can experience a strong dip with increasing Mach number. The characteristics of this dip are sensitive to the nonlinearities in the flow, namely shock-boundary layer interactions.² Particularly for the case of separated flow, accurate numerical predictions of the flutter boundary are very challenging.

The ability to design transonic aircraft to an appropriate degree of flutter conservatism can be compromised by (1) the fact that higher-fidelity flutter prediction may not be included early enough in the design process, and (2) the presence of uncertainties and variability. Regarding point (1), incorporating higher fidelity physics at early design stages can allow transonic flutter mechanisms to be properly accounted for at a stage where modifications are relatively inexpensive, and further ease the burden of expensive wind tunnel and flight flutter testing campaigns.³ Regarding point (2), replacing empirical flutter margins with probabilistic representations (namely, probability of flutter/failure) can help account for deviations in structural and inertial properties within a fleet⁴ and also the fact that transonic flutter can be highly sensitive to uncertainties in the manufactured wing shape.⁵

Taken together, early-stage higher-fidelity probabilistic flutter predictions can help prevent situations where a flutter mechanism is discovered late in the design process.⁶ However, transonic aeroelastic uncertainty quantification (UQ) is expensive, particularly when there is an interest in low-probability, highly-catastrophic events such as flutter, governed by large numbers of random variables.⁷ Examples of higher-fidelity (i.e., computational fluid dynamics, CFD) flutter UQ can be found in Refs. 8-14, with Refs. 13 and 14

*Research Aerospace Engineer, Aeroelasticity Branch, bret.k.stanford@nasa.gov, AIAA Associate Fellow.

†Graduate Student, Department of Statistics.

‡Research Aerospace Engineer, Aeroelasticity Branch, AIAA Member.

§Research Computer Engineer, Durability, Damage Tolerance, and Reliability Branch.

specifically using Reynolds-averaged Navier-Stokes (RANS) solvers. For these solvers, flutter- q is computed with repeated time-integration of the unsteady CFD equations at progressively higher dynamic pressures, until an oscillatory instability is noted. The repeated time integrations, compounded by the multiple function evaluations needed to quantify uncertainties, lead to the high computational cost.

Two cost-reduction schemes are considered here for CFD-based flutter UQ. First, linearized frequency-domain (LFD) methods can be used to compute flutter- q without time integration, instead solving a series of linear equations to assess the aerodynamic response to infinitesimal oscillations of various structural mode shapes.^{15–17} The oscillatory aerodynamic responses (generalized aerodynamic forces, GAFs), can then be used to directly obtain flutter- q via an eigen-equation. Second, gradient-enhanced UQ tools (or in some cases, gradient-enabled) may be considered, given recent progress¹⁸ in computing adjoint derivatives of the LFD-obtained flutter- q with respect to large numbers of parameters. For sampling-based UQ tools, such as polynomial chaos expansion (PCE)¹⁹ or Gaussian processes (GP),⁹ enriching the sampled data with derivatives is a well-known method to reduce sampling costs.^{20,21} For reliability methods such as FORM and SORM (first- and second-order reliability methods),²² derivative information is required for the underlying mechanism of these tools.

This paper demonstrates gradient-enhanced uncertainty quantification and reliability estimation for two aeroelastic problems: the AGARD 445.6 test case²³ and the Common Research Model (CRM).¹³ Aeroelastic computations are conducted with a lower-fidelity linear doublet-lattice method (DLM) aerodynamics,²⁴ or the high-fidelity LFD solver¹⁷ written within a stabilized finite element (SFE) solver²⁵ for the FUN3D software.²⁶ Critically, the outputs from both solvers are the same, namely a complex-valued matrix of generalized aerodynamic forces (GAFs). OpenMDAO²⁷ and the recently-developed Mphys library²⁸ are used to facilitate coupling of these aerodynamic solvers to structural solvers and load/displacement transfer solvers, and further to compute analytical sensitivities of flutter- q with respect to wing shape and sizing variables. Gradient-enhanced UQ and failure probabilities (probability that flutter- q is less than some required threshold) are demonstrated with four different schemes: FORM, SORM, gradient-enhanced PCE, and gradient-enhanced GP with entropy contour location.²⁹

II. Uncertainty Quantification Methods

Random variables are denoted as the vector $\mathbf{x} \in \mathbb{R}^n$; when converted into standard normal space, they are denoted as \mathbf{u} . The function of interest is the scalar $f : \mathbb{R}^n \rightarrow \mathbb{R}$, set to $q_f - q_{limit}$. This function is the difference between the computed flutter point (q_f) and the limiting value (q_{limit} , namely the flight envelope). A failure is denoted by a negative value of f .

A. First-Order Reliability Method

FORM involves finding the most-probable-point (MPP) on the failure surface in the standard normal space, via a gradient-based search:

$$\underset{\mathbf{u}}{\text{minimize}} \quad \sqrt{\mathbf{u}^T \cdot \mathbf{u}} \quad \text{s.t.} : f = 0 \quad (1)$$

If the MPP (i.e., the solution to the optimization problem) is given by \mathbf{u}^* , then the reliability index is $\beta = \sqrt{\mathbf{u}^{*T} \cdot \mathbf{u}^*}$, and the probability of failure is $\mathbb{P}(f < 0) = \Phi(-\beta)$, where Φ is the standard normal cumulative distribution function. This failure probability is exact if the failure surface f is a linear function of \mathbf{u} , and if \mathbf{u} are jointly normal, though errors are introduced otherwise. For this work, the MPP search process stopping criteria occurs when the norm-difference between two consecutive MPP's is 10^{-6} .

B. Second-Order Reliability Method

In the event of nonlinearity in the $f - \mathbf{u}$ relationship, SORM improves on FORM's probability computation by accounting for curvature:

$$\mathbb{P}(f < 0) = \Phi(-\beta) \cdot \prod_{i=1}^{n-1} \frac{1}{\sqrt{1 + \kappa_i \cdot \beta}} \quad (2)$$

where Eq. 2 is the simplest of several formulas that may be used,²² and κ are the principal curvatures of an orthonormal transformation of the Hessian matrix $\nabla_{\mathbf{u}}^2 f$. This Hessian is typically infeasible to compute analytically, but may be estimated from finite differences on the first-order derivatives $df/d\mathbf{u}$ (a scheme that will scale poorly with large numbers of random variables), or also estimated from a quasi-Newton BFGS scheme used to solve Eq. 1.³⁰

C. Gradient-Enhanced Polynomial Chaos Expansion

The PCE scheme used here is a non-intrusive point collocation scheme originally developed in Ref. 19. The output f is assumed to be represented by a weighted combination of basis functions written across the random variable space:

$$\hat{f}(\mathbf{u}) = \sum_{i=0}^P \alpha_i \cdot \Psi_i(\mathbf{u}) \quad (3)$$

where Ψ_i is the i^{th} basis function and α_i is the corresponding weight. For the normal random variables considered here, the basis functions are a total order expansion of Hermite polynomials, comprising $P = (n + s)!/n!/s! - 1$ terms. n is the number of random variables in \mathbf{u} , and s is the desired order of the PCE. Further details regarding the mathematics of this assumed basis may be found in Refs. 19 and 31.

The values of f are sampled at random locations (identified through Latin Hypercube Sampling, LHS) throughout the random variable space, and then α_i terms are computed by solving a least squares problem. If the number of available samples is equal to or greater than $P + 1$, Eq. 3 is an overdetermined least squares problem. Sampling cost grows rapidly with increased n or s , and one way to blunt this growth is to enrich the samples with the gradients $df/d\mathbf{u}$. The matrix equation then becomes:³²

$$\begin{bmatrix} \Psi_0(\mathbf{u}_0) & \Psi_1(\mathbf{u}_0) & \Psi_2(\mathbf{u}_0) & \cdots & \Psi_P(\mathbf{u}_0) \\ \frac{\partial \Psi_0}{\partial \mathbf{u}}(\mathbf{u}_0) & \frac{\partial \Psi_1}{\partial \mathbf{u}}(\mathbf{u}_0) & \frac{\partial \Psi_2}{\partial \mathbf{u}}(\mathbf{u}_0) & \cdots & \frac{\partial \Psi_P}{\partial \mathbf{u}}(\mathbf{u}_0) \\ \Psi_0(\mathbf{u}_1) & \Psi_1(\mathbf{u}_1) & \Psi_2(\mathbf{u}_1) & \cdots & \Psi_P(\mathbf{u}_1) \\ \frac{\partial \Psi_0}{\partial \mathbf{u}}(\mathbf{u}_1) & \frac{\partial \Psi_1}{\partial \mathbf{u}}(\mathbf{u}_1) & \frac{\partial \Psi_2}{\partial \mathbf{u}}(\mathbf{u}_1) & \cdots & \frac{\partial \Psi_P}{\partial \mathbf{u}}(\mathbf{u}_1) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Psi_0(\mathbf{u}_{N_s}) & \Psi_1(\mathbf{u}_{N_s}) & \Psi_2(\mathbf{u}_{N_s}) & \cdots & \Psi_P(\mathbf{u}_{N_s}) \\ \frac{\partial \Psi_0}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) & \frac{\partial \Psi_1}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) & \frac{\partial \Psi_2}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) & \cdots & \frac{\partial \Psi_P}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) \end{bmatrix} \cdot \begin{Bmatrix} \alpha_0 \\ \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_P \end{Bmatrix} = \begin{Bmatrix} f(\mathbf{u}_0) \\ \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_0) \\ f(\mathbf{u}_1) \\ \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_1) \\ \vdots \\ f(\mathbf{u}_{N_s}) \\ \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) \end{Bmatrix} \quad (4)$$

where N_s is the number of samples.

Having solved Eq. 4, the first four statistical moments of f can be computed, and then finally fit to a Pearson model³³ to form a probability density function (PDF). This PDF can then be used to compute the probability of failure, $\mathbb{P}(f < 0)$, by numerically integrating under the curve. Alternatively, failure probabilities can be estimated from Monte Carlo sampling (MCS) of the inexpensive surrogate PCE model itself.

As PCE is an expansion centered about the mean, it may struggle to accurately predict very-low probabilities of failure, particularly if (in the likely event) there are few or no samples generated in the tails of the PDF (i.e., low-probability events). Though work has been conducted to expand PCE methods toward rare events,³⁴ these methods are not explored here.

D. Gradient-Enhanced Gaussian Process with Entropy Contour Location

Similar to the PCE described in the previous section, a GP is a surrogate model that relies on 1) initial samples throughout the random variable space; 2) a regression technique used to tune the hyperparameters of the underlying kernels within a covariance matrix; 3) a final post-processing step where the surrogate model is used to estimate UQ metrics of interest. The main difference from the above PCE is that the GP surrogate intrinsically contains uncertainty information; furthermore a scheme is described below, which ensures that samples are adaptively generated near the failure surface, to better compute low-probability events.

A GP surrogate assumes a prior multivariate normal distribution $f(\mathbf{u}) \sim \mathcal{N}(\mathbf{0}, \Sigma(\mathbf{u}))$ ⁹ where $\Sigma(\mathbf{u})$ is shorthand for $\Sigma(\mathbf{u}, \mathbf{u})$, the $N_s \times N_s$ matrix generated by applying the covariance function $\Sigma(\mathbf{u}_i, \mathbf{u}_j)$ over

all $i, j = 1, \dots, N_s$. Incorporating gradient information requires extending the covariance function to accept gradient input locations.³⁵ With a differentiable kernel (such as the squared exponential), the partial derivatives simply commute: $\Sigma(\mathbf{u}_i, \frac{\partial f}{\partial \mathbf{u}_j}) = \frac{\partial}{\partial \mathbf{u}_j} \Sigma(\mathbf{u}_i, \mathbf{u}_j)$. Similarly, $\Sigma(\frac{\partial f}{\partial \mathbf{u}_i}, \frac{\partial f}{\partial \mathbf{u}_j}) = \frac{\partial^2}{\partial \mathbf{u}_i \partial \mathbf{u}_j} \Sigma(\mathbf{u}_i, \mathbf{u}_j)$. Thus the GP prior for both responses and gradients may be represented as:

$$\left\{ \begin{array}{l} f(\mathbf{u}) \\ \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}) \end{array} \right\} \sim \mathcal{N} \left(\mathbf{0}, \Sigma \left(\left\{ \begin{array}{l} \mathbf{u} \\ \frac{\partial}{\partial \mathbf{u}} \end{array} \right\} \right) \right) \quad \text{where} \quad \Sigma \left(\left\{ \begin{array}{l} \mathbf{u} \\ \frac{\partial}{\partial \mathbf{u}} \end{array} \right\} \right) = \begin{bmatrix} \Sigma(\mathbf{u}, \mathbf{u}) & \Sigma(\mathbf{u}, \frac{\partial f}{\partial \mathbf{u}}) \\ \Sigma(\frac{\partial f}{\partial \mathbf{u}}, \mathbf{u}) & \Sigma(\frac{\partial f}{\partial \mathbf{u}}, \frac{\partial f}{\partial \mathbf{u}}) \end{bmatrix}. \quad (5)$$

Sampling at N_s random locations yields observations $\left\{ \mathbf{u}_{N_s}, f(\mathbf{u}_{N_s}), \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) \right\}$, and finally produces predictions at \mathbf{u}^* with the closed form solution $\hat{f}(\mathbf{u}^*) \sim \mathcal{N}(\mu^*(\mathbf{u}^*), \Sigma^*(\mathbf{u}^*))$, where

$$\mu^*(\mathbf{u}^*) = \Sigma \left(\mathbf{u}^*, \left\{ \begin{array}{l} \mathbf{u}_{N_s} \\ \frac{\partial}{\partial \mathbf{u}_{N_s}} \end{array} \right\} \right) \cdot \Sigma \left(\left\{ \begin{array}{l} \mathbf{u}_{N_s} \\ \frac{\partial}{\partial \mathbf{u}_{N_s}} \end{array} \right\} \right)^{-1} \cdot \left\{ \begin{array}{l} f(\mathbf{u}_{N_s}) \\ \frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_{N_s}) \end{array} \right\} \quad (6)$$

$$\Sigma^*(\mathbf{u}^*) = \Sigma(\mathbf{u}^*) - \Sigma \left(\mathbf{u}^*, \left\{ \begin{array}{l} \mathbf{u}_{N_s} \\ \frac{\partial}{\partial \mathbf{u}_{N_s}} \end{array} \right\} \right) \cdot \Sigma \left(\left\{ \begin{array}{l} \mathbf{u}_{N_s} \\ \frac{\partial}{\partial \mathbf{u}_{N_s}} \end{array} \right\} \right)^{-1} \cdot \Sigma \left(\left\{ \begin{array}{l} \mathbf{u}_{N_s} \\ \frac{\partial}{\partial \mathbf{u}_{N_s}} \end{array} \right\}, \mathbf{u}^* \right). \quad (7)$$

The mean $\mu^*(\mathbf{u}^*)$ may be used as a deterministic prediction of $\hat{f}(\mathbf{u}^*)$, while the covariance $\Sigma^*(\mathbf{u}^*)$ holds point-wise predictive variances that may be utilized for sequential design. The `gpytorch` Python package³⁶ is leveraged to properly handle the required derivative terms across the dimensions of \mathbf{u} .

Having computed the GP predictive distribution, N_{MC} Monte Carlo samples (denoted as \mathbf{u}_{MC}) may be used to estimate probability of failure by averaging this probability across each sample:³⁷

$$\mathbb{P}(f < 0) \approx \mathbb{P}(\hat{f} < 0) \approx \frac{1}{N_{MC}} \cdot \sum_{i=1}^{N_{MC}} \mathbb{P}(\hat{f}(\mathbf{u}_{MC}^{(i)}) < 0) = \frac{1}{N_{MC}} \cdot \sum_{i=1}^{N_{MC}} \Phi \left(\frac{0 - \mu^*(\mathbf{u}_{MC}^{(i)})}{\sqrt{\Sigma^*(\mathbf{u}_{MC}^{(i)})}} \right) \quad (8)$$

where Φ is the standard normal cumulative distribution function.

As noted above, the utility of this failure probability estimate relies on how accurately the GP surrogate can identify the failure region, which is directly tied to the proximity of \mathbf{u}_{N_s} to the failure contour. In the event of small failure probabilities, an LHS (for example) is unlikely to provide samples near the failure boundary. To target this area of the variable space, a sequential design using entropy contour location (ECL)²⁹ is employed. The entropy criterion is given as:

$$\text{ENTROPY}(\mathbf{u} | \hat{f}) = -\mathbb{P}(\hat{f}(\mathbf{u}) \geq 0) \cdot \log(\mathbb{P}(\hat{f}(\mathbf{u}) \geq 0)) - \mathbb{P}(\hat{f}(\mathbf{u}) < 0) \cdot \log(\mathbb{P}(\hat{f}(\mathbf{u}) < 0)). \quad (9)$$

This term measures the degree of uncertainty in failure predictions and is maximized when $\mathbb{P}(\hat{f}(\mathbf{u}) \geq 0) = \mathbb{P}(\hat{f}(\mathbf{u}) < 0) = 0.5$. A gradient search²⁹ is used to locate the area of the variable space that maximizes this entropy term.

A sequential design process can then proceed as follows:

1. Fit a GP surrogate \hat{f} to an initial LHS design of size N_s .
2. Select the next input \mathbf{u}_{N_s+1} at the location that maximizes entropy.
3. Evaluate the response $f(\mathbf{u}_{N_s+1})$ and the derivative $\frac{\partial f}{\partial \mathbf{u}}(\mathbf{u}_{N_s+1})$.
4. Update the GP surrogate \hat{f} with this new data.
5. Repeat steps 2-4 as needed.

Sequential acquisitions will focus learning near the failure contour, leading to more accurate failure probability estimates with fewer model evaluations. Though not investigated here, a batch sequential design strategy may be utilized by identifying more than one new entropy-maximizing design per cycle, and evaluating the response of each design in parallel.

For relatively large failure probabilities, the cost of Eq. 8 is not prohibitive given the small expense of the GP surrogate. However, for lower probabilities the computational cost of gradient-enhanced GP predictions

can still become a bottleneck. Gradient observations multiply the sample size by $n+1$, consequently inflating the cubic costs of GP inference. GP predictions without derivative observations are much faster to compute. To leverage this speed, a multi-fidelity importance sampling scheme (MFIS)³⁸ is utilized, where GP without gradients is the lower-fidelity model and the gradient-enhanced GP is the high-fidelity model. This allows for efficient computation of smaller failure probabilities.

III. Aeroelastic Modeling

As noted above, aeroelastic computations are conducted with either lower-fidelity DLM-based aerodynamics²⁴ or the high-fidelity LFD-based aerodynamics.¹⁷ For either model, an OpenMDAO-centric²⁷ approach has been utilized to efficiently compute aeroelastic response and derivative information, namely via the Mphys library.²⁸ Mphys has been developed as a collaborative effort seeking to standardize high-fidelity multiphysics problems (in this case, unsteady aeroelastic problems) in OpenMDAO. The Mphys library also provides utilities to help set up these multiphysics models, which can quickly grow in complexity. A key feature of this library is swappable fidelity, provided that both fidelity analyses can match the same standardized interfaces. The LFD tool can be replaced with a lower-fidelity DLM solver without modifying other aspects of the problem. Similarly, the structural model or load/displacement transfer scheme may be swapped without modifying the aerodynamic solver setup.

A simplified view of the Mphys-generated aeroelastic workflow can be seen in the extended design structure matrix (XDSM) diagram³⁹ of Fig. 1. Random variables are broken into aerodynamic shape variables (\mathbf{u}_{shape}) and structural sizing variables (\mathbf{u}_{sizing}). The wing geometry discipline defines the structural mesh (\mathbf{x}_{struct}) and the aerodynamic mesh (\mathbf{x}_{aero}). A steady aerodynamic solver is then used to compute the steady background flow over the wing. For the CFD-based workflow, this would involve driving a steady flow residual to 0, via the FUN3D/SFE tool,²⁵ to compute the flow solution \mathbf{q}_0 . For a DLM-based workflow, this entire step can be omitted, since the flutter outputs of Fig. 1 are independent of the steady flow, for linear physics.

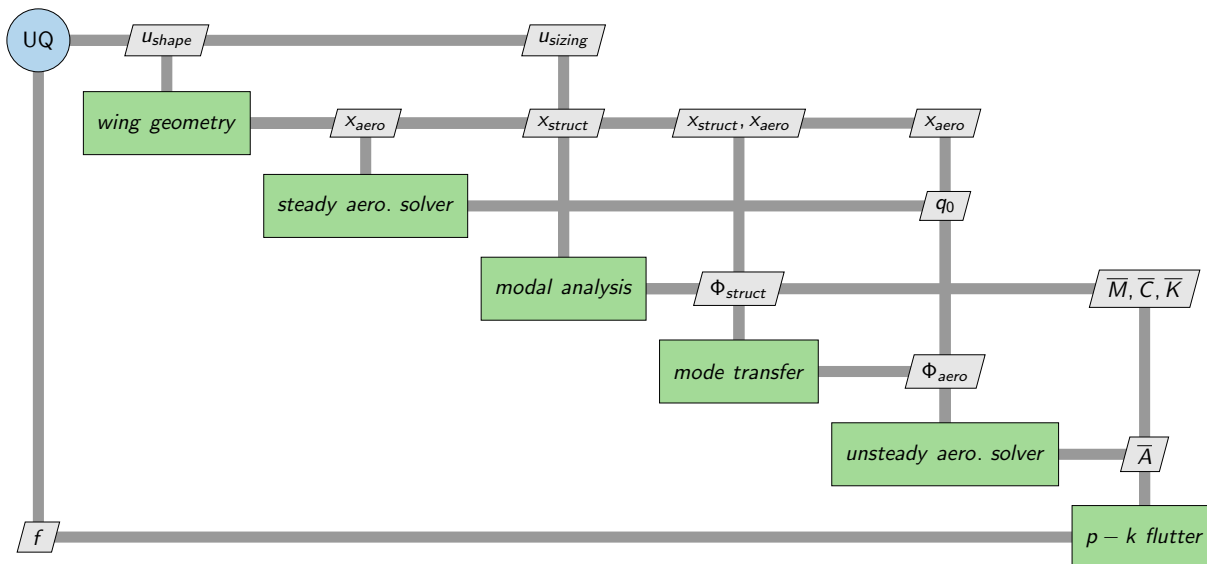


Figure 1: XDSM diagram of the aeroelastic flutter UQ process.

Next, a structural modal analysis discipline defines the stiffness (\mathbf{K}) and mass (\mathbf{M}) matrices, then solves the free-vibration eigen-problem for the vibration modes:

$$(\mathbf{K} - \omega_j^2 \cdot \mathbf{M}) \cdot \phi_j = \mathbf{0} \quad (10)$$

where ω_j is the j^{th} natural frequency, and ϕ_j is the corresponding mode shape. The first J mode shapes may be grouped into the matrix Φ_{struct} , where J is far smaller than the number of degrees of freedom in the finite element analysis. The mode shapes are normalized as:

$$\phi_j^T \cdot \mathbf{M} \cdot \phi_j = 1 \quad (11)$$

The modally-reduced stiffness matrix is defined as $\bar{\mathbf{K}} = \Phi_{struct}^T \cdot \mathbf{K} \cdot \Phi_{struct}$, with similar definitions for mass and damping matrices. The structural mode shapes Φ_{struct} must then be transferred from the structural mesh (\mathbf{x}_{struct}) to the aerodynamic surface mesh (\mathbf{x}_{aero}) to form the term Φ_{aero} . This can be done with either a simple radial-basis function interpolation scheme, or the weighted least squares scheme from Ref. 40.

For the CFD-based workflow, the unsteady aerodynamic solver module involves solving the following equations:

$$\mathbf{M}(\mathbf{q}, \mathbf{x}_G, \dot{\mathbf{x}}_G) \cdot \dot{\mathbf{q}} + \mathbf{R}(\mathbf{q}, \mathbf{x}_G, \dot{\mathbf{x}}_G) = \mathbf{0}, \quad (12)$$

where \mathbf{q} is the flow state vector, \mathbf{x}_G is the vector of volume mesh coordinates, \mathbf{R} is the spatial residual, and \mathbf{M} is the mass matrix. One can then assume small oscillatory perturbations of the form $\mathbf{x}_G = \mathbf{x}_{G_0} + \sum_{j=1}^n \hat{\mathbf{x}}_{G_j} \cdot e^{i \cdot \omega \cdot t}$, and $\mathbf{q} = \mathbf{q}_0 + \sum_{j=1}^n \hat{\mathbf{q}}_j \cdot e^{i \cdot \omega \cdot t}$, where (for example) \mathbf{q}_0 is a steady-state solution, $\hat{\mathbf{q}}$ is a Fourier coefficient, and ω is the oscillation frequency. Each Fourier coefficient corresponds to a structural mode shape of interest. Using these assumptions, the LFD equations can be written as:

$$\left(i \cdot \omega \cdot \mathbf{M} \Big|_{\mathbf{x}_{G_0}, \mathbf{q}_0} + \frac{\partial \mathbf{R}}{\partial \mathbf{q}} \Big|_{\mathbf{x}_{G_0}, \mathbf{q}_0} \right) \cdot \hat{\mathbf{q}}_j + \left(i \cdot \omega \cdot \frac{\partial \mathbf{R}}{\partial \dot{\mathbf{x}}_G} \Big|_{\mathbf{x}_{G_0}, \mathbf{q}_0} + \frac{\partial \mathbf{R}}{\partial \mathbf{x}_G} \Big|_{\mathbf{x}_{G_0}, \mathbf{q}_0} \right) \cdot \hat{\mathbf{x}}_{G_j} = \mathbf{0}. \quad (13)$$

Equation 13 represents a linear complex-valued system of equations, which must be solved once for each structural mode shape j (namely via the surface motion Φ_{aero} , which is used to compute the oscillatory volume mesh motion $\hat{\mathbf{x}}_{G_j}$) and each frequency value ω of interest. The complex-valued oscillatory flow fields $\hat{\mathbf{q}}_j$ can be used to compute a set of generalized aerodynamic forces (GAFs), $\bar{\mathbf{A}}$.

For the DLM-based workflow, the unsteady aerodynamic solver module involves solving the following complex-valued equations at a range of pre-selected frequencies:

$$\bar{\mathbf{A}}(i \cdot k) = \mathbf{Q} \cdot \mathbf{D}(i \cdot k)^{-1} \cdot \left(\frac{\partial \Phi_{aero}}{\partial \mathbf{x}} + \frac{i \cdot k}{b} \cdot \Phi_{aero} \right) \quad (14)$$

where b is the reference length, \mathbf{D} is the DLM aerodynamic influence coefficient, \mathbf{Q} is a matrix that converts aerodynamic pressures into modal forces, and $\bar{\mathbf{A}}$ are again the GAFs. The reduced frequency is defined as $k = \omega \cdot b / U$, where U is the flow speed.

Flutter- q is finally computed with the p - k equations:

$$\left(\frac{U^2}{b^2} \cdot p^2 \cdot \bar{\mathbf{M}} + \frac{U}{b} \cdot p \cdot \bar{\mathbf{C}} + \bar{\mathbf{K}} - q \cdot \bar{\mathbf{A}}(i \cdot k) \right) \cdot \mathbf{v} = \mathbf{0} \quad (15)$$

where \mathbf{v} is the aeroelastic eigenvector, and p is the nondimensional Laplace parameter. The real and imaginary parts of p are expressed as $p = g + i \cdot k$, where g is a damping parameter. The exact flutter point q_f is defined as the dynamic pressure at which g becomes zero, and can be computed by driving the following residual to 0:

$$\mathbf{R}_f(\mathbf{x}_f) = \left\{ \begin{array}{c} \text{Re}(p) \\ k - \text{Im}(p) \end{array} \right\} \quad (16)$$

The corresponding solution vector used to solve Eq. 16 is:

$$\mathbf{x}_f = \left\{ \begin{array}{c} q \\ k \end{array} \right\} \quad (17)$$

It is important to note that the result in Fig. 1 is only valid for non-lifting flows; for lifting flows with a static aeroelastic component, a substantially more complex OpenMDAO model is required (see Ref. 41) with a nonlinear block Gauss-Seidel (NLBGS) scheme needed to iterate between the static structural solver and the steady flow solver.

Finally, by properly computing the analytical partial derivatives local to the inputs and outputs of each discipline in Fig. 1, OpenMDAO is able to compute the final derivative of the flutter output (q_f , and then $f = q_f - q_{limit}$) with respect to the random inputs (\mathbf{u}), needed to conduct the UQ analyses outlined above. The derivatives of the unsteady LFD-based aerodynamic component are particularly challenging to compute, with details presented in Refs. 18 and 41.

IV. Results

Two computational aeroelastic test cases are considered here: the AGARD 445.6 test case²³ and the CRM.¹³ The former test case is a very-thin wing at 0.0° angle of attack (AoA); even for transonic Mach numbers, the nonlinearities in the flow are weak. For this case then, differences between FUN3D/LFD and DLM will be minor. The CRM, on the other hand, is a higher-aspect ratio swept wing composed of conventional supercritical airfoils. Complex shock structures will appear over the wing in the transonic flow regime, which DLM will be unable to predict.

A. AGARD Test Case

This test case is a sidewall-mounted wing with an aspect ratio of 1.65, a taper ratio of 0.66, a sweep angle of 45°, a semispan of 2.5 feet, and a root chord of 1.833 feet. Four independent random variables are identified for the AGARD 445.6 test case:

- An additive variable that changes the wing sweep: $\mathcal{N}(0.0, 1.0)$
- An additive variable that changes the wing span: $\mathcal{N}(0.0, 1.0)$
- A multiplicative variable that scales all of the structural shell thickness values: $\mathcal{N}(1.0, 0.02)$
- A multiplicative variable that scales the material orientation of the mahogany structure: $\mathcal{N}(1.0, 0.05)$

1. DLM-Based Results

Using a 3rd-order gradient-enhanced PCE (Eq. 4), q_f PDFs at various subsonic Mach numbers are shown on the left of Fig. 2, as predicted by DLM. Sobol indices (global sensitivity parameters analytically computed from the PCE coefficients α_i)⁴² are shown on the right of Fig. 2, for each random variable. The nonlinearity in these PDFs is generally weak; at Mach 0.95, for example, the skewness (third statistical moment) is 0.454 and the kurtosis (fourth statistical moment) is 3.418.




Figure 2: Probabilistic flutter boundaries of the AGARD 445.6 test case, as computed with DLM.

At Mach 0.95, failure probabilities may be computed for a range of limiting q_{limit} values; results are shown in Fig. 3 for 250,000 MCS samples that directly query the OpenMDAO flutter model, for comparison with FORM, SORM, gradient-enhanced PCE ($N_s=11$ samples), and gradient-enhanced GP with ECL ($N_s=5$

initial samples, with 15 sequential ECL samples). At the lowest q_{limit} value considered (0.3 psi), too-few failures in the 250,000 MCS sample set prevented any failure probability estimation. Similarly, GP-based failures at this lowest level were computed with the MFIS scheme noted above;³⁸ at larger q_{limit} values, MCS samples on the GP model itself were sufficient (Eq. 8).

Across all of Fig. 3, the failure probabilities computed with SORM and GP/ECL generally match well, with SORM slightly closer to MCS at the lowest q_{limit} value for which MCS is available (0.35 psi). As noted above, Hessian matrices needed for SORM are estimated from a quasi-Newton procedure.³⁰ Because there are only 4 random variables for this case, the Hessian can also be more-accurately computed from finite difference steps on the gradient df/du . At the lowest q_{limit} , the SORM \mathbb{P}_f value for this finite difference Hessian is $1.4886 \cdot 10^{-7}$, in reasonable agreement with the quasi-Newton value ($1.5601 \cdot 10^{-7}$). FORM overpredicts the MCS failure probabilities across all of Fig. 3, indicative of some curvature in the failure boundary. Not surprisingly, PCE is the least accurate at the lowest q_{limit} (relative to the SORM/GP values), with a $\mathbb{P}_f = 6.7907 \cdot 10^{-7}$: this scheme does not use any samples near the failure boundary (unlike the other methods), and so failure probabilities are computed by extrapolating information from mean-centered statistical moments.




Figure 3: DLM-based failure/flutter probabilities of the AGARD 445.6 test case, as computed from MCS, FORM, SORM, PCE, and GP, at Mach 0.95.

The UQ tools driven by sampling (PCE, GP/ECL) will provide probability estimates that depend on the initial LHS sampling. For GP, the spread in computed failure probability is shown in Fig. 4 as computed from 10 different random LHS initial designs. MCS is utilized as a reference in this figure along with SORM, because: 1) MCS is unable to provide reliability estimates at the lowest limit value, and 2) Fig. 3 shows that MCS and SORM are generally in good agreement. As above, there are five initial samples (N_s), followed by 15 steps of the ECL scheme. The figure shows substantial drops in the variability of \mathbb{P}_f with the addition of gradient-enhancement (which is leveraged in both the contour-locating design and final reliability calculations). To preserve a direct comparison between methods, the MFIS scheme for the lowest q_{limit} was replicated without gradients by simply using the typical GP as both the high-fidelity and low-fidelity models.

Similar information is shown in Fig. 5 for the case of gradient-enriched PCEs, again using 10 different LHS designs of 11 samples each. Data for a PCE without gradients is not shown here, since gradient-enriched PCEs are, in this work, viewed as a means to obtain similar results to a PCE without gradients, though with far fewer samples (as will be discussed below). The results in Fig. 5 largely echo the PCE results in Fig. 3, with high variability in the low- \mathbb{P}_f area of the curve, and an overprediction relative to SORM and MCS, as well. Variability in computing large failure probabilities is much better, however, owing to the mean-center nature of the PCE.

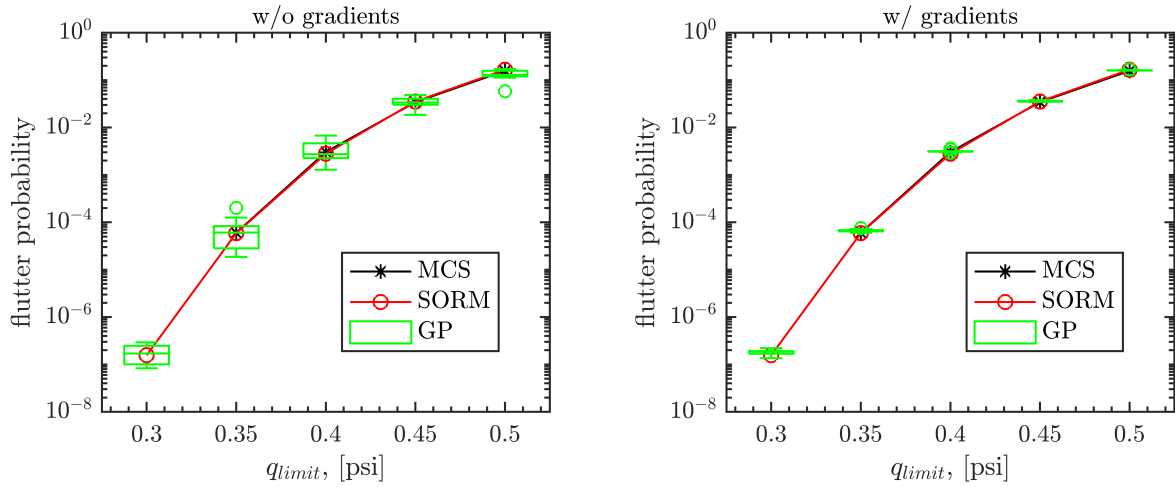


Figure 4: DLM-based failure/flutter probabilities of the AGARD 445.6 test case, as computed from GP with and without gradient information included.




Figure 5: DLM-based failure/flutter probabilities of the AGARD 445.6 test case, as computed from PCE with gradient information included.

2. CFD-Based Results

Next, the DLM-based solver is replaced with an inviscid CFD-based solver (Euler flow), with a focus on flutter boundaries at Mach 0.5. The volume mesh used in FUN3D contains 145,000 nodes, and the total cost of computing flutter boundaries is roughly two orders of magnitude higher than the DLM-based workflow. Given the increased computational cost of the flutter predictions, greater attention is paid to the relative cost of each of the UQ methods, in addition to their accuracy. First, PCE-based results are shown in Fig. 6 and Table 1, in terms of the PDFs, Sobol indices, and statistical moments; as noted, this PDF is very similar to the DLM-based PDF shown in Fig. 2 at Mach 0.5, owing to the weak flow nonlinearities for this configuration. The PCE without gradients utilized 70 LHS samples: for a 3rd-order expansion ($s=3$) and 4 random variables ($n=4$), at least $(n+s)!/n!/s!=35$ samples are needed to avoid an underdetermined least squares problem, but Ref. 19 recommends an oversampling ratio of 2. No effort is made here to thoroughly demonstrate the minimum number of samples needed for gradient-enriched PCE, but Fig. 2 shows adequate comparisons with the 70-sample PCE when using 11 gradient-enriched samples (oversampling=0.33). For

this AGARD test case, the cost of computing flutter- q and its derivatives is roughly double the cost of computing flutter- q alone. Increasing the number of random variables will improve this cost benefit: the number of required samples will grow (for both standard PCE and gradient-enriched PCE), but the cost of computing derivatives will not, given the adjoint-based tools used here.

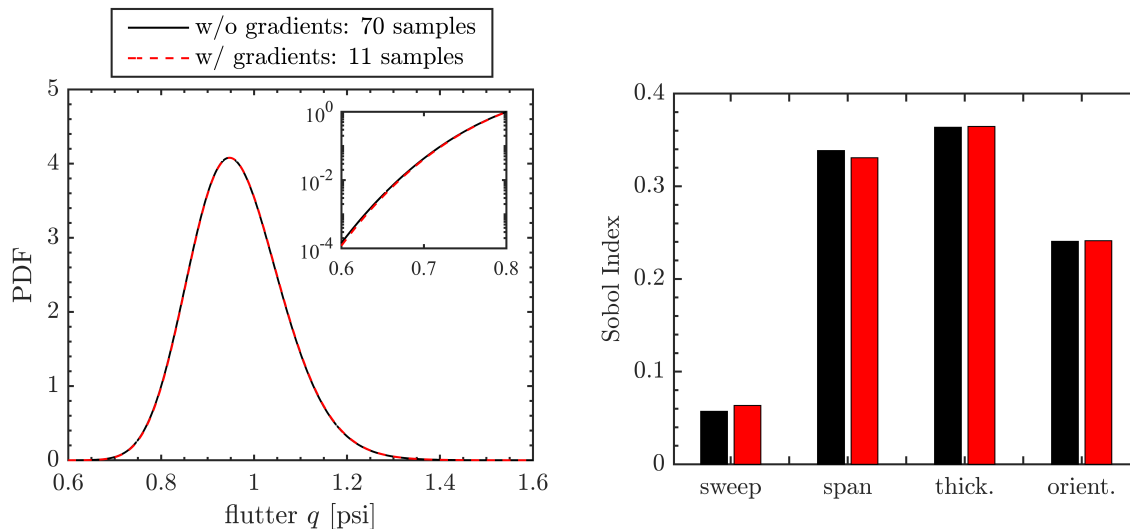


Figure 6: CFD-based PDFs and Sobol indices of the AGARD 445.6 test case at Mach 0.5.

Table 1: CFD-based statistical moments for the AGARD 445.6 test case at Mach 0.5.

	PCE	mean	standard deviation	skew	kurtosis	$\mathbb{P}_f (q_{limit}=0.75 \text{ psi})$
w/o gradients: 70 samples		0.96531 psf	0.10033 psf	0.38976	3.32289	0.00749
w/gradients: 11 samples		0.96593 psf	0.10035 psf	0.39376	3.32453	0.00722

Figure 7 demonstrates flutter probability convergence for FORM, SORM, and GP/ECL, for a q_{limit} value of 0.6 psi. The left part of this figure shows the convergence of the flutter failure boundary; for FORM and SORM, this value starts at 0.949 psi, which is the deterministic flutter- q value at this Mach number, and is very similar to the mean value listed in Table 1. The FORM/SORM results converges to the q_{limit} (i.e., the MPP) within 4 iterations, at which point the flutter probability is $1.599 \cdot 10^{-6}$ (FORM) and $1.002 \cdot 10^{-6}$ (SORM, which augments the FORM probabilities with Hessian information, Eq. 2). For GP/ECL, the first five data points on the left of Fig. 7 are the initial LHS sample. The ECL search (which is a global search, rather than the local search process of FORM and SORM) begins at iteration 5, and samples 5 additional points near the q_{limit} boundary. The final flutter probability for GP/ECL is $5.236 \cdot 10^{-7}$, somewhat lower than predicted by SORM. The “true” probability for this case is unknown, because the CFD-based model is too expensive for a MCS simulation, as opposed to the less-expensive DLM-based models in the previous section.

Deeper insight into the uncertain aeroelastic physics is shown in the next few figures, in terms of results for the MPP design found by FORM/SORM, as compared with the baseline deterministic result. The designs themselves are shown in Fig 8, where the color contour indicates the structural shell thickness (inches), and the thin black lines indicates material orientation vectors. In addition to decreasing the shell thickness and slightly rotating the material orientation vectors further into the flow, the MPP also increases the wing span and decreases the wing sweep. Taken together, these 4 design changes (corresponding to the 4 random variables) decrease the flutter- q from 0.949 to 0.6 psi, as seen in Fig. 9 where an eigenvalue with a positive real part induces a flutter instability.

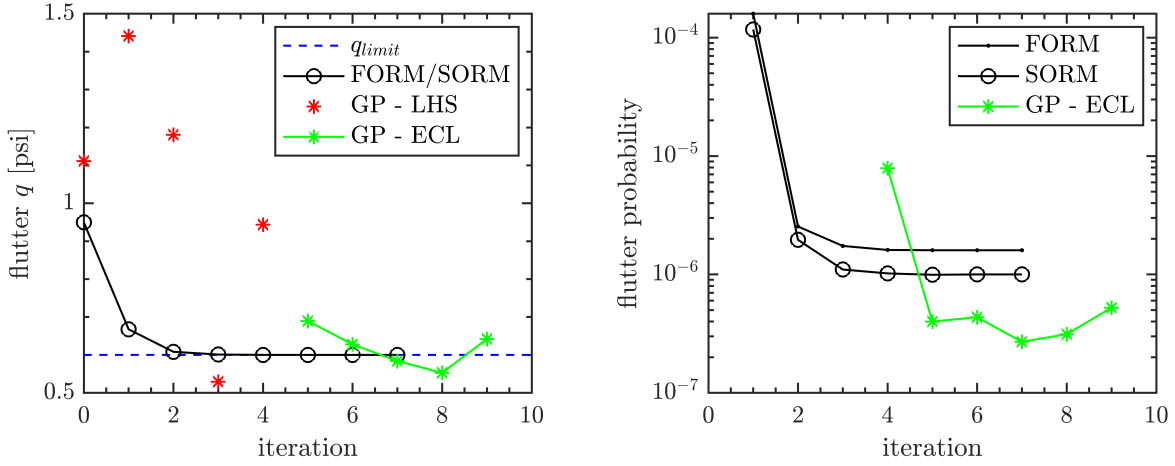


Figure 7: CFD-based flutter probability convergence of the AGARD 445.6 test case at Mach 0.5.




Figure 8: Baseline deterministic design (left) and the MPP design (right), for the AGARD 445.6 test case at Mach 0.5: shell thickness contours (inches) and material orientation vectors (thin black lines).

B. Common Research Model

This test case is a generic transport configuration jig shape, developed in Ref. 43, with a wing span of 58.7 m, a mean aerodynamic chord of 7.0 m, an aspect ratio of 9, a taper ratio of 0.275, and a sweep angle of 35° . The aluminum wingbox structure consists of an upper skin, a lower skin, a leading edge spar (located at 10% chord at the root and 35% at the tip), a trailing edge spar (60% at both the root and the tip), and 43 ribs oriented perpendicular to the leading edge. All shell members (ribs, spars, skins) are outfitted with T-shaped stiffeners, where the stiffeners are not modeled explicitly, but instead smeared into the shell stiffness properties. The structure along the leading and trailing edges of the wing, along with the nacelle and pylon, is not explicitly modeled, though an inertial effect is captured with a series of lumped masses attached to the wing box via interpolation elements. Additional modeling details can be found in Ref. 13.

Six random variables are considered for this configuration: five that change the structural thickness of the wingbox (inboard skin thickness, outboard skin thickness, inboard spar thickness, outboard spar thickness, rib thickness), and one that changes the orientation of the skin stiffeners. Each random variable is a multiplicative variable that scales the baseline structural properties, and is considered to be a Gaussian normal variable defined as $\mathcal{N}(1.0, 0.05)$. Only inviscid CFD-based results are shown here, as opposed to DLM. The results in this section are more preliminary than the previous section: only PCE-based results are shown, and the static aeroelastic load redistribution (likely important for this transonic lifting flow) is ignored here.




Figure 9: CFD-based eigenvalue migration of the AGARD 445.6 test case at Mach 0.5.

PCE results are shown in Fig. 10 for the CRM at Mach 0.6 and Mach 0.85 (0.0° AoA). PDFs of the CFD-computed flutter- q are shown in the upper right, via a 3rd-order PCE with 168 samples: 84 samples are required to prevent an under-determined system, and the over-sampling factor is set to 2. Alternatively, a gradient-enhanced PCE can reduce the number of samples required by a factor of ten, with little discernible shift in the resulting PDFs. Global sensitivity indices are shown in the upper right of Fig. 10, with the stiffer orientation variable (which will have a large effect on the bend-twist coupling of the wing) showing the largest impact on the flutter boundary.

Transonic flow is clearly evident in the CFD result at Mach 0.85 in Fig. 10, and Fig. 11 expands on this idea further by varying the AoA at the same Mach number. A relationship between flutter- q and AoA is an inherently nonlinear one: the linear DLM flow solver is incapable of predicting any relationship at all. Fig. 11 shows the expected² rise in flutter- q with AoA, and the PCE modeling shows substantial increases in flutter variability with increased AoA also. Via the Sobol indices, the randomness in stiffer orientation becomes progressively more important with increased AoA, at the expense of the five structural thickness variables.

V. Conclusions

This paper has demonstrated the use of four different UQ techniques for computing reliability analyses of dynamic aeroelastic systems, with a focus on low-probability flutter events. The four UQ methods are: 1) first order reliability method; 2) second order reliability method; 3) polynomial chaos expansion; 4) Gaussian processes with entropy contour location. Aeroelastic flutter is computed via OpenMDAO/Mphys modeling, which helps enable analytical gradients of flutter behavior with respect to parametric inputs. Those sensitivities can then in turn be used to enhance the UQ computations: the two reliability methods (FORM, SORM) depend on a gradient search to compute failure probabilities, while the two sampling methods (PCE, GP) will be able to lower sampling costs with the addition of gradients.

Two test cases are utilized here: the AGARD 445.6 and the Common Research Model, with the latter's results more preliminary in nature. For the AGARD case, of the four UQ methods, SORM produces results with the fewest required samples, and (when MCS results are available for comparison) the highest accuracy. However, the expectation is that for problems with stronger nonlinearities in the failure boundaries, or problems with multi-modality in the PDF (from competing flutter mechanisms, see Ref. 44), the GP/ECL method will likely be superior in terms of its ability to accurately compute low probabilities of failure within a limited computational budget. Furthermore, the ECL process can be leveraged in a parallel computing environment by adding multiple points at once along the failure contour, whereas SORM is inherently sequential. The CRM test case under viscous transonic flows, including the static aeroelastic dependency (which was ignored here) will likely be a good test case to verify these claims, and will be the focus of future work.

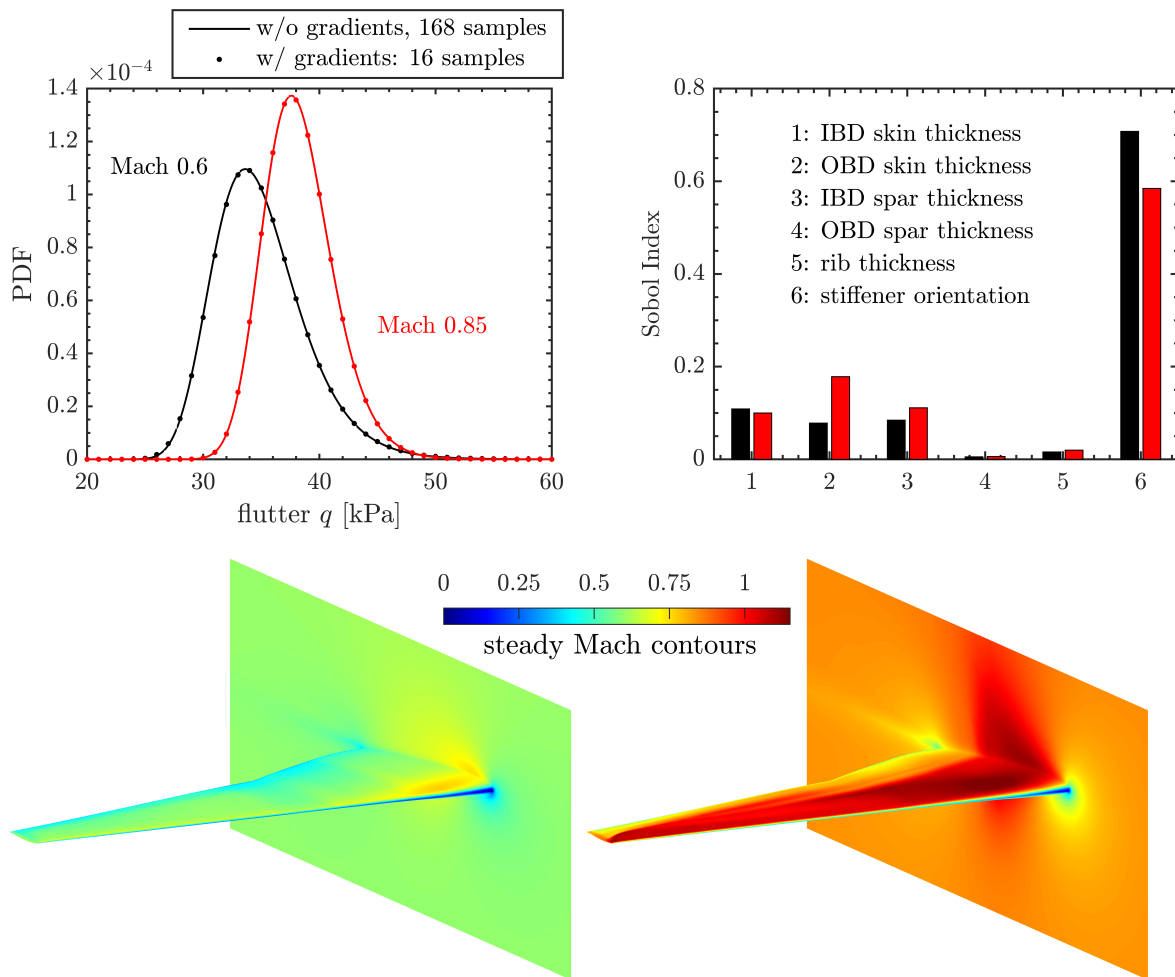


Figure 10: PCE modeling of an uncertain CRM flutter boundary at Mach 0.6 and 0.85 (0° AoA): probability density functions (top left); global sensitivity indices (top right); steady Mach field contours at Mach 0.6 (bottom left) and 0.85 (bottom right).

Acknowledgments

This work is supported by the Transformational Tools and Technologies (TTT) project of the NASA Transformative Aeronautics Concepts Program (TACP). Computational resources for this work are provided by the NASA Langley K cluster.

References

- ¹Pettit, C., "Uncertainty Quantification in Aeroelasticity: Recent Results and Research Challenges," *Journal of Aircraft*, Vol. 41, No. 5, pp. 1217-1229, 2004.
- ²Edwards, J., "Calculated Viscous and Scale Effects on Transonic Aeroelasticity," *Journal of Aircraft*, Vol. 45, No. 6, pp. 1863-1871, 2008.
- ³Cooper, J., "Towards Faster and Safer Flight Flutter Testing," RTO AVT Symposium on Reduction of Military Vehicle Acquisition Time and Cost Through Advanced Modeling and Virtual Simulation, NATO Paper 089, March 2003.
- ⁴Pitt, D., Haudrich, D., Thomas, M., Griffin, K., "Probabilistic Aeroelastic Analysis and Its Implications on Flutter Margin Requirements," *AIAA Structures, Structural Dynamics, and Materials Conference*, Schaumburg, IL, April 2008.
- ⁵Housman, J., Kiris, C., "Overset Grid Simulations for the Second AIAA Aeroelastic Prediction Workshop," *AIAA SciTech Forum*, Grapevine, TX, January 2017.
- ⁶Gates, D., "In Person: Fitzgerald's Fix for the Boeing 747-8 Earns Aviation Honors," *Seattle Times*, <http://www.seattletimes.com/business/in-person-fitzgeralds-fix-for-boeing-747-8-earns-aviation-honors> [retrieved 2021].
- ⁷Beran, P., Stanford, B., Schrock, C., "Uncertainty Quantification in Aeroelasticity," *Annual Review of Fluid Mechanics*, Vol. 49, pp. 361-386, 2017.

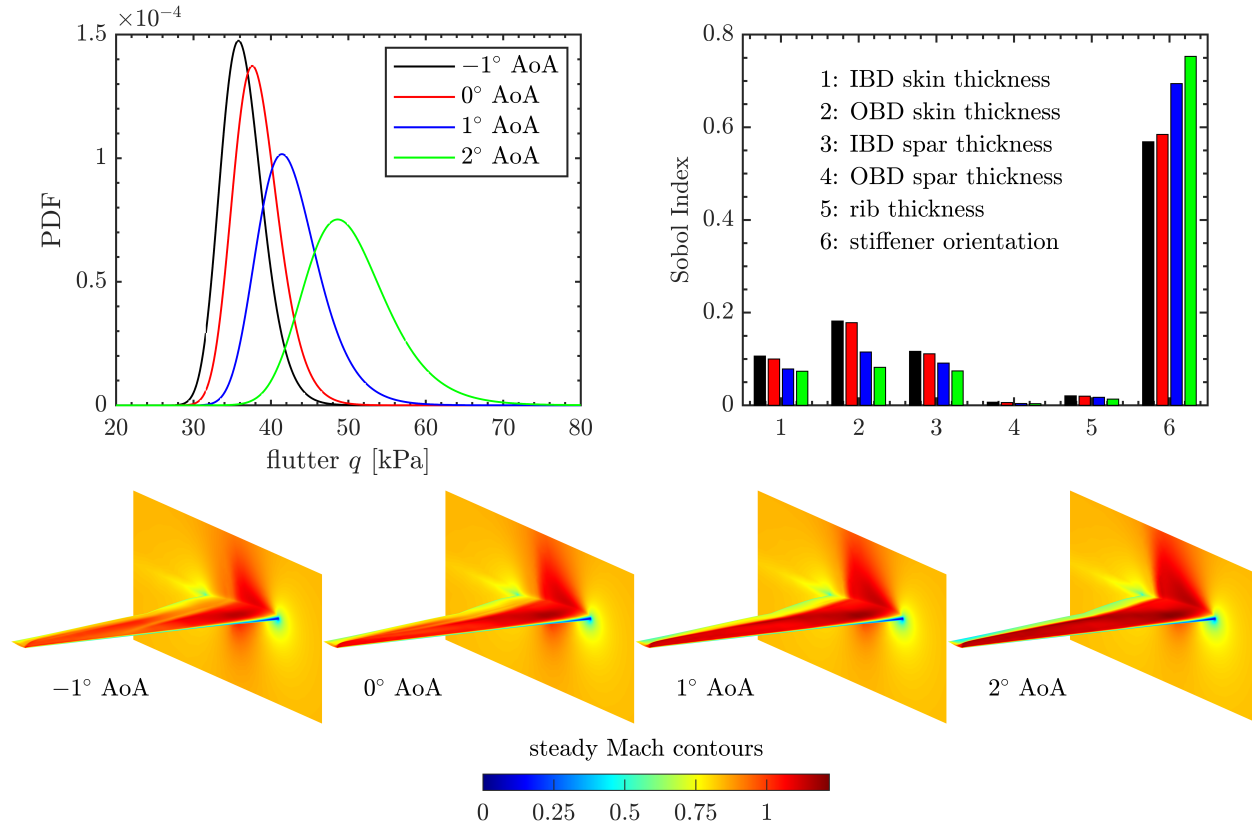


Figure 11: PCE modeling of an uncertain CRM flutter boundary at Mach 0.85 (-1° , 0° , 1° , 2° AoA): probability density functions (top left); global sensitivity indices (top right); steady Mach field contours (bottom).

⁸Hosder, S., Walters, R., Balch, M., “Efficient Uncertainty Quantification Applied to the Aeroelastic Analysis of a Transonic Wing,” *AIAA Aerospace Sciences Meetings and Exhibit*, Reno, NV, January 7-10, 2008.

⁹Rasmussen, C., Williams, C., *Gaussian Processes for Machine Learning*, MIT Press, Cambridge, MA, 2005.

¹⁰Witteveen, J., Bijl, H., “Effect of Randomness on Multi-Frequency Aeroelastic Responses Resolved by Unsteady Adaptive Stochastic Finite Elements,” *Journal of Computational Physics*, Vol. 228, No. 18, pp. 7025-2045, 2009.

¹¹Marques, S., Badcock, K., Khodaparast, H., Mottershead, J., “Transonic Aeroelastic Stability Predictions under the Influence of Structural Variability,” *Journal of Aircraft*, Vol. 47, No. 4, pp. 1229-1239, 2010.

¹²Nikbay, M., Kuru, M., “Reliability Based Multidisciplinary Optimization of Aeroelastic Systems with Structural and Aerodynamic Uncertainties,” *Journal of Aircraft*, Vol. 50, No. 3, pp. 708-715, 2013.

¹³Stanford, B., Massey, S., “Uncertainty Quantification of the FUN3D-Predicted NASA CRM Flutter Boundary,” *AIAA SciTech Forum*, Grapevine, TX, January 2017.

¹⁴Righi, M., Greco, P., Da Ronch, A., “Uncertainty Quantification of CFD-Based Flutter Prediction,” *AIAA SciTech Forum*, January 2021.

¹⁵Timme, S., Badcock, K., “Transonic Aeroelastic Instability Searches using Sampling and Aerodynamic Model Hierarchy,” *AIAA Journal*, Vol. 49, No. 6, pp. 1191-1201, 2012.

¹⁶Thormass, R., Widhalm, M., “Linear-Frequency-Domain Predictions of Dynamic-Response Data for Viscous Transonic Flows,” *AIAA Journal*, Vol. 51, No. 11, pp. 2540-2557, 2013.

¹⁷Jacobson, K., Stanford, B., Wood, S., Anderson, W., “Frequency-Domain Flutter Analysis and Sensitivities with Stabilized Finite Elements,” *AIAA SciTech Forum*, Orlando, FL, January 6-10, 2020.

¹⁸Jacobson, K., Stanford, B., Wood, S., Anderson, W., “Adjoint-based Sensitivities of Flutter Predictions based on the Linearized Frequency-Domain Approach,” *AIAA SciTech Conference*, January 2021.

¹⁹Hosder, S., Walters, R., Balch, M., “Point-Collocation Nonintrusive Polynomial Chaos Method for Stochastic Computational Fluid Dynamics,” *AIAA Journal*, Vol. 48, No. 12, pp. 2721-2730, 2010.

²⁰Roderick, O., Anitescu, M., Fischer, P., “Polynomial Regression Approaches Using Derivative Information for Uncertainty Quantification,” *Nuclear Science and Engineering*, Vol. 164, pp. 122-139, 2010.

²¹Alekseev, A., Navon, I., Zelentsov, M., “The Estimation of Functional Uncertainty using Polynomial Chaos and Adjoint Equations,” *International Journal for Numerical Methods in Fluids*, Vol. 67, No. 3, pp. 328-341, 2011.

- ²²Rackwitz, R., "Reliability Analysis - A Review and Some Perspectives," *Structural Safety*, Vol. 23, No. 4, pp. 365-395, 2001.
- ²³Yates, C., "AGARD Standard Aeroelastic Configurations for Dynamic Response. Candidate Configuration I - Wing 445.6," NASA TM-100492, August, 1987.
- ²⁴Blair, M., "A Compilation of the Mathematics Leading to the Doublet Lattice Method," WL-TR-92-3028, August, 1992.
- ²⁵Anderson, W., Newman, J., Karman, S., "Stabilized Finite Elements in FUN3D," *Journal of Aircraft*, Vol. 55, No. 2, pp. 696-714, 2017.
- ²⁶Biedron, R., Carlson, J., Derlaga, J., Gnoffo, P., Hammond, D., Jacobson, K., Jones, W., Kleb, B., Lee-Rausch, E., Nielsen, E., Park, M., Rumsey, C., Thomas, J., Thompson, K., Walden, A., Wang, L., Wood, W., "FUN3D Manual: 13.7," NASA TM-2020-5010139.
- ²⁷Gray, J., Hwang, J., Martins, J., Moore, K., Naylor, B., "OpenMDAO: An Open-Source Framework for Multidisciplinary Design, Analysis, and Optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, No. 4, pp. 1075-1104, 2019.
- ²⁸Yildirim, A., Jacobson, K., Stanford, B., Gray, J., Mader, C., Martins, J., Kennedy, G., "A Modular Multiphysics Simulation Framework Using OpenMDAO," *AIAA Aviation Forum*, virtual event, June 15-19, 2020.
- ²⁹Cole, D., Gramacy, R., Warner, J., Bomarito, G., Leser, P., Leser, W., "Entropy-Based Adaptive Design for Contour Finding and Estimating Reliability," *arXiv preprint*, arXiv:2105.11357, 2021.
- ³⁰Pericaro, G., Santos, S., Ribeiro, A., Matioli, L., "HLRF-BFGS Optimization Algorithm for Structural Reliability" *Applied Mathematical Modeling*, Vol. 39, No. 7, pp. 2025-2035, 2015.
- ³¹Eldred, M., "Recent Advances in Non-Intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Analysis and Design," AIAA Paper 2009-2274.
- ³²Rallabhandi, S., West, T., Nielsen, E., "Uncertainty Analysis and Robust Design of Low-Boom Concepts Using Atmospheric Adjoints," *Journal of Aircraft*, Vol. 54, No. 3, pp. 902-917, 2016.
- ³³Johnson, N., *Continuous Univariate Distributions*, Wiley-Interscience, 1994.
- ³⁴Schobi, R., Sudret, B., Marelli, S., "Rare Event Estimation using Polynomial-Chaos Kriging," *Journal of Risk and Uncertainty in Engineering Systems*, Vol. 3, No. 2, D4016002, 2017.
- ³⁵Solak, E., Murray-Smith, R., Leithead, W., "Derivative Observations in Gaussian Process Models of Dynamic Systems," in *Advances in Neural Information Processing Systems*, MIT Press, Cambridge, MA, 2002.
- ³⁶Gardner, J., Pleiss, G., Bindel, D., Weinberger, K., Wilson, A., "GPpyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration," in *Advances in Neural Information Processing Systems*, MIT Press, Cambridge, MA, 2002.
- ³⁷Picheny, V., Ginsbourger, D., Roustant, O., Haftka, R., Kim, N. "Adaptive Designs of Experiments for Accurate Approximation of a Target Region," *Journal of Mechanical Design*, Vol. 132, No. 7, 2010.
- ³⁸Peherstorfer, B., Cui, T., Marzouk, Y., Willcox, K., "Multifidelity Importance Sampling," *Computer Methods in Applied Mechanics and Engineering*, Vol. 300, pp. 490-509, 2016.
- ³⁹Lambe, A., Martins, J., "Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes" *Structural and Multidisciplinary Optimization*, Vol. 46, pp. 273-284, 2012.
- ⁴⁰Kiviaho, J., Kennedy, G., "Efficient and Robust Load and Displacement Transfer Scheme using Weighted Least Squares" *AIAA Journal*, Vol. 57, No. 5., pp. 2238-2243, 2019.
- ⁴¹Jacobson, K., Stanford, B., "Flutter-Constrained Optimization with the Linearized Frequency-Domain Approach" *AIAA Scitech Forum*, 2022, to appear.
- ⁴²Sudret, B., "Global sensitivity analysis using polynomial chaos expansion," *Reliability Engineering and System Safety*, Vol. 93, No. 7, 2008, pp. 964-979.
- ⁴³Brooks, T., Kenway, G., Martins, J., "Benchmark Aerostructural Models for the Study of Transonic Aircraft Wings," *AIAA Journal*, Vol. 56, No. 7, pp. 2840-2855, 2018.
- ⁴⁴Scarth, C., Sartor, P., Cooper, J., Weaver, P., Silva, G., "Robust and Reliability-Based Aeroelastic Design of Composite Plate Wings," *AIAA Journal*, Vol. 55, No. 10, pp. 3540-3552, 2017.