

# A multifidelity method for locating aeroelastic flutter boundaries

Alexandre N. Marques\*, Max M. J. Opgenoord†, Remi R. Lam‡, and Anirban Chaudhuri§  
*Massachusetts Institute of Technology, Cambridge, MA, 02139*

Karen E. Willcox¶  
*University of Texas at Austin, Austin, TX 78712*

This paper introduces a multifidelity method that can produce accurate estimates of the flutter boundary at a reduced cost by combining information from low- and high-fidelity aeroelastic models. Estimating the flutter boundary in the presence of non-linear aerodynamic phenomena is challenging because high-fidelity aeroelastic models are expensive to evaluate, and flutter analysis requires many model evaluations. On the other hand, relatively inexpensive approximate aeroelastic models (low-fidelity models) exist and are routinely applied to reduce the cost of estimating flutter, albeit with lower accuracy. The multifidelity method introduced here uses an active learning algorithm to leverage information from low-fidelity models. A multifidelity statistical surrogate is used to fit damping coefficient estimates computed with different aeroelastic models. This surrogate is used to estimate the uncertainty in the prediction of the flutter boundary, which drives the selection of new evaluations. The effectiveness of the multifidelity method is demonstrated by estimating the aeroelastic flutter boundary of a typical section model at a cost 85% lower when compared to the bisection method. Four aeroelastic models are considered in this example: three models (including the high-fidelity model) use a computational fluid dynamics solver based on the Euler equations, whereas one model uses a two-dimensional doublet-lattice method.

## Nomenclature

$\mathcal{A}$	=	operator that models convective and diffusive phenomena governing the airflow dynamics
$b$	=	semi-chord length
$c_\ell$	=	lift coefficient
$c_m$	=	momentum coefficient with respect to the elastic axis

---

\*Postdoctoral Associate, Department of Aeronautics and Astronautics, AIAA Member.

†Former Postdoctoral Associate, Department of Aeronautics and Astronautics, AIAA Member.

‡Former Postdoctoral Associate, Department of Aeronautics and Astronautics. Now at DeepMind. AIAA Member.

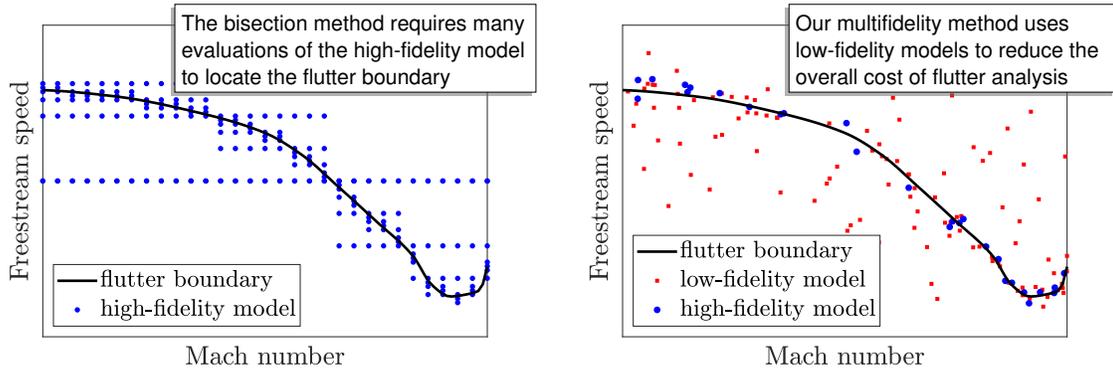
§Postdoctoral Associate, Department of Aeronautics and Astronautics, AIAA Member.

¶Professor, Oden Institute for Computational Engineering and Sciences, AIAA Fellow.

- $d_j$  = damping of aeroelastic mode  $j$
- $f$  = Gaussian process surrogate model
- $F_a$  = aerodynamic forces
- $F_{el}$  = elastic forces
- $F_b$  = body forces, including effects of rigid body acceleration
- $K$  = stiffness matrix
- $M$  = mass matrix
- $m$  = mean function in Gaussian process surrogate model
- $M_\infty$  = Mach number at freestream
- $p$  = cost of evaluating aeroelastic model
- $q_\infty$  = dynamic pressure at freestream
- $r_\theta$  = radius of gyration around the elastic axis
- $V_\infty$  = speed at freestream
- $V_\mu$  = speed index,  $V_\infty/\sqrt{\mu}b\omega_\theta$
- $\mathbf{x}$  = vector of parameters that affect aeroelastic stability
- $x_{cg}$  = position of center of gravity in semi-chords, measured from mid-chord
- $x_{ea}$  = position of elastic axis in semi-chords, measured from mid-chord
- $y$  = quantity of interest modeled by statistical surrogate model,  $\tanh(s\gamma)/s$  for some constant  $s$
- $w$  = mass of airfoil per unit span
- $\gamma$  = aeroelastic damping coefficient
- $\eta$  = degrees-of-freedom of structure
- $\eta_a$  = degrees-of-freedom of airflow
- $\mu$  = mass ratio,  $\mu = w/\pi\rho_\infty b^2$
- $\Sigma$  = covariance kernel in Gaussian process surrogate model
- $\omega_j$  = angular frequency of aeroelastic mode  $j$
- $\omega_\theta$  = uncoupled natural angular frequency of pitch mode in vacuum, in rad/s
- $\omega_h$  = uncoupled natural angular frequency of heave mode in vacuum, in rad/s
- $\rho_\infty$  = density at freestream
- $\theta$  = pitch angle, in radians
- $\xi$  = heave displacement, in semi-chords

## I. Introduction

Accurate flutter prediction for modern aircraft must take into consideration non-linear aerodynamic effects. High-fidelity aeroelastic models based on computational fluid dynamics (CFD) can predict such non-linear behavior [1–11], but remain too expensive for engineering design where typically hundreds to thousands of model evaluations are required per design iteration. To reduce the cost of design, several low-fidelity models have been developed for flutter analysis over the years, such as the doublet-lattice method (DLM) [12–14], but may not capture the full physics of the system. This paper introduces a multifidelity method that leverages data from multiple aeroelastic models to predict flutter boundaries with accuracy comparable to that of a high-fidelity model, while reducing the overall evaluation cost. The proposed method uses principles from information theory to quantify the amount of information each model can offer about the location of the flutter boundary, and an active learning technique to decide how to select model evaluations, accounting for both their accuracy and their computational cost. The result is a multifidelity method that uses many evaluations of inexpensive low-fidelity models and only a few carefully selected evaluations of the high-fidelity model, as illustrated in Fig. 1. As our results will show, this method can be significantly more efficient than using the high-fidelity model only with a standard bisection method.



**Fig. 1** Illustrative comparison between the bisection method and our multifidelity method to locate the aeroelastic flutter boundary.

Our multifidelity method is based on CLoVER [15], an active learning algorithm that combines information from multiple sources to locate function contours. CLoVER is an iterative algorithm that adaptively acquires new data points (evaluations of aeroelastic models). At every iteration, CLoVER fits a multifidelity statistical surrogate model to data available from all aeroelastic models. The surrogate model is then used to quantify the uncertainty in the prediction of the flutter boundary, and also used as a generative model to estimate the one-step lookahead reduction in uncertainty. CLoVER then evaluates the aeroelastic model expected to lead to the largest reduction in uncertainty per unit cost. An open source implementation of CLoVER is available at <https://github.com/anmarques/CLoVER>.

Other multifidelity methods for flutter analysis have been proposed. Dribusch et al. [16] applied a multifidelity

adaptive support vector machine (SVM) algorithm to estimate the aeroelastic flutter boundary of an airfoil subject to structural non-linearities. Timme and Badcock [17] use cokriging to model the entries of the Schur complement matrix of the aeroelastic system leveraging data from aerodynamic models ranging from full potential to Reynolds-Averaged Navier-Stokes equations. Thelen and Leifsson [18] use a similar cokriging technique to model the entries of the aerodynamic influence coefficient matrix, combining data from DLM and the Euler equations. Both cokriging approaches show that accurate flutter predictions of two-degrees-of-freedom airfoils are possible with few evaluations of the high-fidelity model per structural mode.

There are important distinctions between the method proposed here and other multifidelity methods. Our method accommodates analyses based on multiple aeroelastic models and does not require any hierarchy between the models, except for designating one of the models as the high-fidelity model. Although it may be relatively straightforward to find a hierarchy between physical models, aeroelastic models can reflect different combinations of physical models, discretizations, and numerical parameters which do not necessarily result in a clear hierarchy. In addition, Timme and Badcock [17] introduce sampling strategies that resemble active learning techniques. However, their sampling strategy is only applied to the highest fidelity model, whereas our active learning algorithm decides which model to evaluate at each iteration.

Another approach to incorporate high-fidelity data into aeroelastic analysis is producing data-driven low-fidelity aeroelastic models. Several techniques have been used for this purpose. Some techniques are based on measuring the aerodynamic response to small perturbations of the structure using high-fidelity simulations, and fitting the results with linearized models [19–25]. Other techniques use principles of model order reduction to identify a lower-dimensional model of the aeroelastic response, using methods such as proper orthogonal decomposition [26–30], harmonic balance [31], and Volterra series [26, 32]. In principle, any of these data-driven low-fidelity models can be applied with the multifidelity method proposed here.

The remainder of this paper is organized as follows. In Sec. II, we present a formal definition of the aeroelastic flutter boundary. In Sec. III, we discuss the details of the multifidelity method. We apply this method to locate the aeroelastic flutter boundary of a typical section problem, and the results are in Sec. IV. In Sec. V, we discuss several topics that can affect the performance of the multifidelity method. Finally, in Sec. VI, we present conclusions and list opportunities for future work.

## II. Aeroelastic flutter

The behavior of an aeroelastic system is governed by the dynamics of the structure under the influence of aerodynamic forces. Here we represent the state of the structure and airflow by discrete sets of degrees-of-freedom. The vector of generalized coordinates (translations/rotations, modal coordinates)  $\boldsymbol{\eta} \in \mathbb{R}^N$  defines the state of the structure, while the vector of flow variables (velocity, gas properties)  $\boldsymbol{\eta}_a \in \mathbb{R}^{N_a}$  defines the state of the airflow. The dynamics of a general

aeroelastic system are governed by

$$M\dot{\boldsymbol{\eta}} + F_{\text{el}}(\boldsymbol{\eta}) = F_{\text{a}}(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}, \ddot{\boldsymbol{\eta}}, \boldsymbol{\eta}_{\text{a}}, \dot{\boldsymbol{\eta}}_{\text{a}}, q_{\infty}) + F_{\text{b}}, \quad (1\text{a})$$

$$\dot{\boldsymbol{\eta}}_{\text{a}} = \mathcal{A}(M_{\infty}, \text{Re}, \boldsymbol{\eta}_{\text{a}}, \boldsymbol{\eta}, \dot{\boldsymbol{\eta}}), \quad (1\text{b})$$

where  $M$  is the mass matrix,  $F_{\text{el}}$  denotes elastic forces,  $F_{\text{b}}$  denotes body forces,  $F_{\text{a}}$  denotes aerodynamic forces acting on the structure, and  $\mathcal{A}$  represents convective and diffusive phenomena that govern the evolution of the airflow—e.g., (1b) can represent the Navier-Stokes equations. In general, the airflow dynamics depend on the dynamic pressure,  $q_{\infty}$ , Mach number,  $M_{\infty}$ , and the Reynolds number,  $\text{Re}$ , all measured in freestream conditions. In addition, when the structure is not restrained, a non-inertial coordinate system aligned with the structure's center of gravity is employed and the effects of rigid body accelerations are included in  $F_{\text{b}}$ .

Let  $\{\boldsymbol{\eta}_0, \boldsymbol{\eta}_{\text{a}0}\}$  denote a state of aeroelastic equilibrium, i.e.,

$$F_{\text{el}}(\boldsymbol{\eta}_0) = F_{\text{a}}(\boldsymbol{\eta}_0, \boldsymbol{\eta}_{\text{a}0}, q_{\infty}) + F_{\text{b}},$$

$$\mathcal{A}(M_{\infty}, \text{Re}, \boldsymbol{\eta}_{\text{a}0}, \boldsymbol{\eta}_0) = 0.$$

Aeroelastic flutter is characterized by the response of the aeroelastic system (1), initially at equilibrium, to an infinitesimal disturbance  $\delta\boldsymbol{\eta}$  in the generalized coordinates.\* In the infinitesimal limit, the response of the structure can be approximated as linear with respect to the disturbance magnitude:

$$\boldsymbol{\eta}(t) = \boldsymbol{\eta}_0 + \|\delta\boldsymbol{\eta}\| \sum_{j=1}^N \bar{\boldsymbol{\eta}}_j e^{(d_j + i\omega_j)t}, \quad (2)$$

where  $\bar{\boldsymbol{\eta}}_j$  denotes the  $j$ -th mode of vibration of the aeroelastic system (out of a total of  $N$  modes), with damping  $d_j$  and angular frequency  $\omega_j$ . Let  $J$  denote the vibration mode with largest damping, i.e.,  $d_J = \max_j d_j$ . This mode determines the stability of the aeroelastic system:  $d_J = 0$  indicates the onset of aeroelastic flutter, and the system is said to flutter if  $d_J > 0$ . Aeroelastic damping is also commonly described in terms of the damping coefficient:  $\gamma = d_J/\omega_J$ .

Let  $\boldsymbol{x}$  denote the vector of parameters that affect the stability of the aeroelastic system (e.g., Mach number, speed, mass, stiffness, etc.). The *aeroelastic flutter boundary* is defined as the set of conditions

$$\mathcal{Z} = \{\boldsymbol{x} \mid \gamma(\boldsymbol{x}) = 0\}. \quad (3)$$

---

\*Because this definition is based on an infinitesimal disturbance around aeroelastic equilibrium, it excludes limit-cycle oscillations.

### III. A multifidelity method for aeroelastic flutter analysis

The multifidelity method proposed here is based on CLoVER (Contour Location Via Entropy Reduction) [15], an active learning algorithm that combines information from multiple models to locate the zero contour of an expensive-to-evaluate function. In the case of aeroelastic flutter analysis, one wants to locate the flutter boundary  $\mathcal{Z}$ , which corresponds to the zero contour of the aeroelastic damping coefficient  $\gamma(\mathbf{x})$ . In general, high-fidelity evaluations of  $\gamma$  are expensive. Our method leverages information from low-fidelity aeroelastic models to estimate the flutter boundary accurately at a reasonable cost. In Sec. III.A, we explain what constitutes an aeroelastic model, and establish the notation for the remainder of this section. In Sec. III.B, we introduce a quantity of interest that approximates  $\gamma$  in the vicinity of the flutter boundary and is less sensitive to numerical errors in other regions of the parameter space.

The multifidelity method has three main ingredients:

- A multifidelity statistical surrogate model that fits data from low- and high-fidelity aeroelastic models, encoding correlations between the different fidelities. We discuss the surrogate model in Sec. III.C.
- A measure of uncertainty about the location of the flutter boundary estimated by the statistical surrogate model, as detailed in Sec. III.D.
- A decision mechanism that selects which model and aeroelastic condition to evaluate at each iteration such that the uncertainty about the location of the flutter boundary is reduced the most, per unit computational cost. We discuss this mechanism in Sec. III.E.

Finally, in Sec. III.F, we show how these ingredients are combined to form the multifidelity method.

#### A. Aeroelastic models and notation

In this paper, the term *aeroelastic model* is used to designate the combination of modeling assumptions, numerical discretizations, and numerical algorithms used to:

- (i) describe the operators  $F_{el}$ ,  $F_a$ ,  $F_b$ , and  $\mathcal{A}$ ,
- (ii) solve the aeroelastic problem (1), and
- (iii) estimate the aeroelastic damping coefficient, defined by (2).

We assume that we have access to  $N_m$  aeroelastic models, each of which provides an estimate of the damping coefficient  $\gamma$ . We denote the estimate from model  $\ell$  as  $\gamma_\ell$ , made with computational cost  $p_\ell$ ,  $\ell \in \{0, \dots, N_m - 1\}$ . In this paper we measure computational cost as the ratio of CPU times, taking the high-fidelity model as reference (i.e.,  $p_0 = 1$ ).

Let  $\ell = 0$  denote the aeroelastic model that results in our most accurate (and likely most expensive) estimate of aeroelastic damping coefficient. We refer to aeroelastic model  $\ell = 0$  as the *high-fidelity model*, and aim to estimate the aeroelastic flutter boundary defined by

$$\mathcal{Z}_0 = \{\mathbf{x} \mid \gamma_0(\mathbf{x}) = 0\}. \quad (4)$$

The aeroelastic models  $\ell = 1, \dots, N_m - 1$  are referred to as *low-fidelity models*.

## B. Quantity of interest

Estimating the aeroelastic damping coefficient  $\gamma$  when  $|\gamma| \gg 0$  is a challenge. Large absolute values of  $\gamma$  correspond to fast decay or rapid increase of oscillations, which may lead to situations where numerical simulations offer little data to make accurate estimates of damping and frequency. Although we are only interested in accurate estimates of  $\gamma$  in the vicinity of the flutter boundary, inaccurate estimates elsewhere may lead to spurious and sharp variations in the surrogate model, which hinders the performance of the approach proposed here.

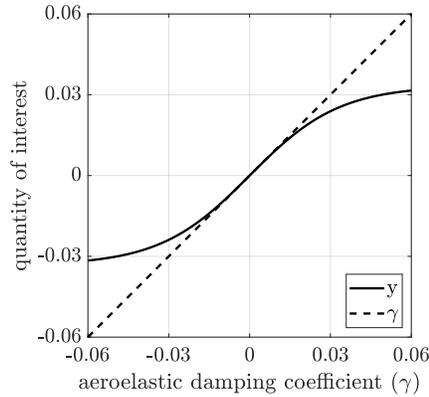
We remedy this issue by defining an alternate quantity of interest that is less sensitive to errors when  $|\gamma| \gg 0$ , which is given by

$$y = \frac{1}{s} \tanh(s\gamma). \quad (5)$$

The parameter  $s$  controls the threshold above which  $y$  becomes insensitive to errors. Here we set  $s = 30.0$ . Note that

- $\gamma = 0 \Leftrightarrow y = 0$  (the zero contour of  $y$  coincides with the aeroelastic flutter boundary),
- $y \approx \gamma$  for  $|\gamma| \ll 1/s$ , and
- $dy/d\gamma \approx 0$  for  $|\gamma| > 1/s$ .

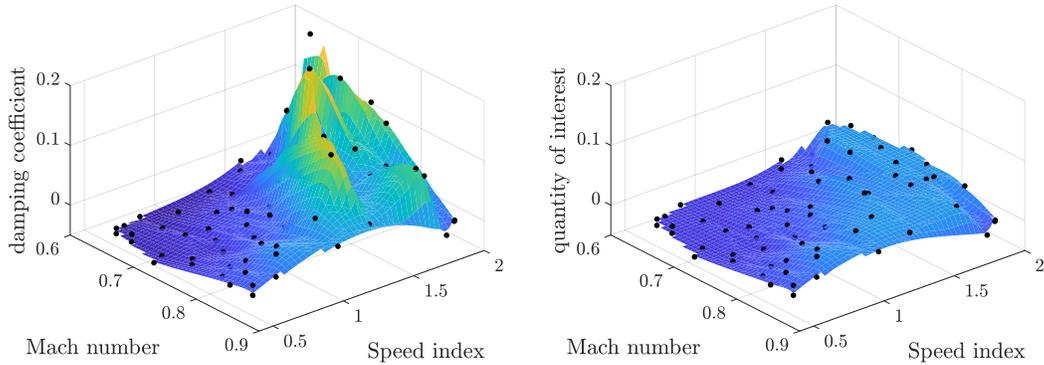
Figure 2 illustrates the relationship between the quantity of interest  $y$  and the aeroelastic damping coefficient  $\gamma$ .



**Fig. 2** Quantity of interest  $y = (1/s) \tanh(s\gamma)$  for  $s = 30.0$ .  $y$  approximates  $\gamma$  in the vicinity of  $\gamma = 0$ .

Figure 3 contrasts the variation of aeroelastic damping coefficient and the quantity of interest defined in (5) for the flutter problem described in Sec. IV. One can observe that the quantity of interest varies more smoothly over the parameter space, facilitating the construction of a surrogate model that approximates the data accurately.

Another approach to avoid the issue of estimating damping coefficient is treating aeroelastic flutter as a classification problem [16, 33]. In this approach one only needs to identify whether the aeroelastic system is stable or unstable, which is simpler than computing estimates of damping coefficient. On the other hand, damping coefficient estimates provide



**Fig. 3** Aeroelastic damping coefficient (left) and the quantity of interest defined in (5) (right) for the typical section problem investigated in Sec. IV. The quantity of interest has the same zero contour as the aeroelastic damping coefficient but varies more smoothly. Surfaces obtained by cubic interpolation of data points shown as black dots.

the advantages of offering a measure of distance to the flutter boundary, and enabling multifidelity frameworks based on model discrepancy (either directly or through a regularized form, such as the quantity of interest defined in (5)).

### C. Multifidelity surrogate model

Our method is influenced by work on multi-information source optimization [34, 35] and uses the statistical surrogate model introduced by Poloczek et al. [35]. This model constructs a single Gaussian process (GP) surrogate that simultaneously approximates the low- and high-fidelity models, and thus exploits relationships between them.

Let  $f$  denote the surrogate model, with  $f(\ell, \mathbf{x})$  being the GP that represents the belief about function  $y_\ell(\mathbf{x})$ ,  $\ell = 0, \dots, N_m - 1$ . The construction of the surrogate follows from two modeling choices:

- (i) a GP approximation to  $y_0(\mathbf{x})$  denoted by  $f(0, \mathbf{x})$ , i.e.,  $f(0, \mathbf{x}) \sim GP(m_0, \Sigma_0)$ , and
- (ii) independent GP approximations to the discrepancies  $\delta_\ell(\mathbf{x}) = y_\ell(\mathbf{x}) - y_0(\mathbf{x})$ , i.e.,  $\delta_\ell \sim GP(m_\ell, \Sigma_\ell)$  with  $\text{Cov}(\delta_\ell(\mathbf{x}), f(0, \mathbf{x}')) = 0$  and  $\text{Cov}(\delta_\ell(\mathbf{x}), \delta_{\ell'}(\mathbf{x}')) = \mathbb{1}_{\ell, \ell'} \Sigma_\ell(\mathbf{x}, \mathbf{x}')$ , where  $\mathbb{1}_{\ell, \ell'}$  denotes the Kronecker's delta.

In the definitions above,  $m_\ell$  denotes the prior mean function and  $\Sigma_\ell$  denotes the covariance kernel of the corresponding GP. As a consequence of these modeling choices, the surrogate model  $f(\ell, \mathbf{x}) = f(0, \mathbf{x}) + (1 - \mathbb{1}_{\ell, 0})\delta_\ell(\mathbf{x})$  is also a GP,

$f \sim GP(m, \Sigma)$ , with

$$\begin{aligned}
m(\ell, \mathbf{x}) &= \mathbb{E}[f(\ell, \mathbf{x})] \\
&= \mathbb{E}[f(0, \mathbf{x})] + (1 - \mathbb{1}_{\ell,0})\mathbb{E}[\delta_{\ell}(\mathbf{x})] \\
&= m_0(\mathbf{x}) + (1 - \mathbb{1}_{\ell,0})m_{\ell}(\mathbf{x}),
\end{aligned} \tag{6}$$

$$\begin{aligned}
\Sigma((\ell, \mathbf{x}), (\ell', \mathbf{x}')) &= \text{Cov}(f(\ell, \mathbf{x}), f(\ell', \mathbf{x}')) \\
&= \text{Cov}\left((f(0, \mathbf{x}) + (1 - \mathbb{1}_{\ell,0})\delta_{\ell}(\mathbf{x})), (f(0, \mathbf{x}') + (1 - \mathbb{1}_{\ell,0})\delta_{\ell}(\mathbf{x}'))\right) \\
&= \text{Cov}(f(0, \mathbf{x}), f(0, \mathbf{x}')) + (1 - \mathbb{1}_{\ell,0}) \text{Cov}(\delta_{\ell}(\mathbf{x}), \delta_{\ell'}(\mathbf{x}')) \\
&\quad + (1 - \mathbb{1}_{\ell,0}) \text{Cov}(f(0, \mathbf{x}), \delta_{\ell'}(\mathbf{x}')) + (1 - \mathbb{1}_{\ell,0}) \text{Cov}(f(0, \mathbf{x}'), \delta_{\ell}(\mathbf{x})) \\
&= \Sigma_0(\mathbf{x}, \mathbf{x}') + (1 - \mathbb{1}_{\ell,0})\mathbb{1}_{\ell, \ell'}\Sigma_{\ell}(\mathbf{x}, \mathbf{x}').
\end{aligned} \tag{7}$$

Note that the multifidelity surrogate model  $f(\ell, \mathbf{x})$  is a standard GP with a particular form of mean function, Eq. (6), and covariance kernel, Eq. (7). Therefore, assimilating data follows from standard tools of Gaussian process regression [36]. Consider  $n$  samples evaluated at  $X_n = \{(\ell^i, \mathbf{x}^i)\}_{i=1}^n$ , which result in observations  $Y_n = \{y_{\ell^i}(\mathbf{x}^i)\}_{i=1}^n$ . We denote the posterior GP of  $f$ , conditioned on  $\{X_n, Y_n\}$ , as  $f^n$ , with mean  $m^n$  and covariance kernel  $\Sigma^n$ .

In practice,  $m_{\ell}$  and  $\Sigma_{\ell}$  are selected from one of the standard parameterized classes of mean functions and covariance kernels [36]. Selecting the parameters of these functions and kernels (known as hyperparameters) is an important consideration. To mitigate difficulties associated with estimating hyperparameters with small amounts of data, the multifidelity method updates its estimate of hyperparameters whenever the high-fidelity model is evaluated. When the high-fidelity model is evaluated, the multifidelity method also evaluates all low-fidelity models at the same location in parameter space. The resulting data are then used to estimate the hyperparameters of the high-fidelity surrogate and of the model discrepancy surrogates independently using maximum likelihood estimates [36].

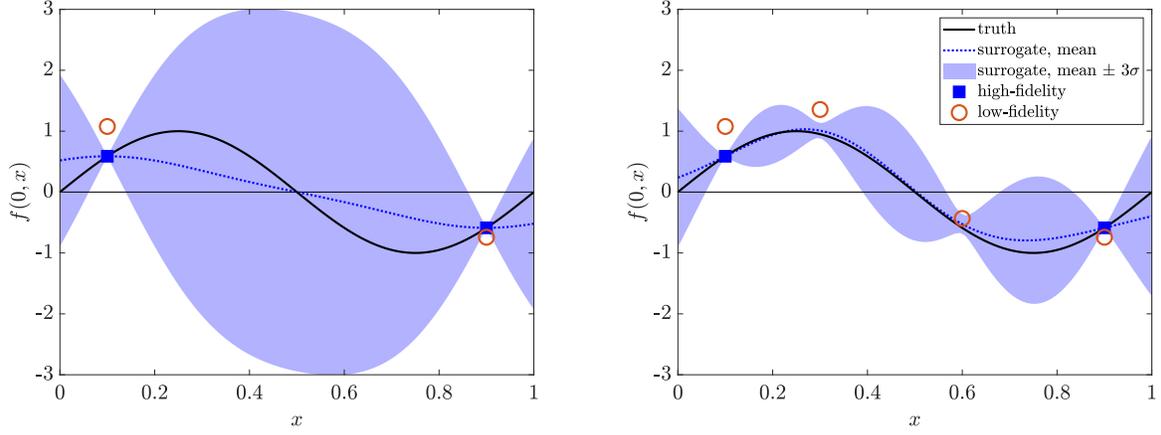
One important consequence of the surrogate construction described above is that low-fidelity data affect the surrogate representation of the high-fidelity model. Figure 4 illustrates this effect.

#### D. Flutter boundary uncertainty

The uncertainty about the location of the flutter boundary is measured by applying the concept of *contour entropy* [15] to the surrogate model discussed above. Contour entropy measures the uncertainty of the zero contour estimated by a statistical surrogate model by defining a discrete random variable associated with point-wise predictions and integrating the associated entropy.<sup>†</sup>

---

<sup>†</sup>In information theory, entropy is a measure of the uncertainty in the outcome of a random process [37].



**Fig. 4** Low-fidelity data improves prediction of high-fidelity function. **Left:** two observations of low- and high-fidelity models. Multifidelity surrogate fits high-fidelity data and learns correlation between different models. **Right:** two additional observations of the low-fidelity model result in significant improvement of surrogate representation of high-fidelity model.

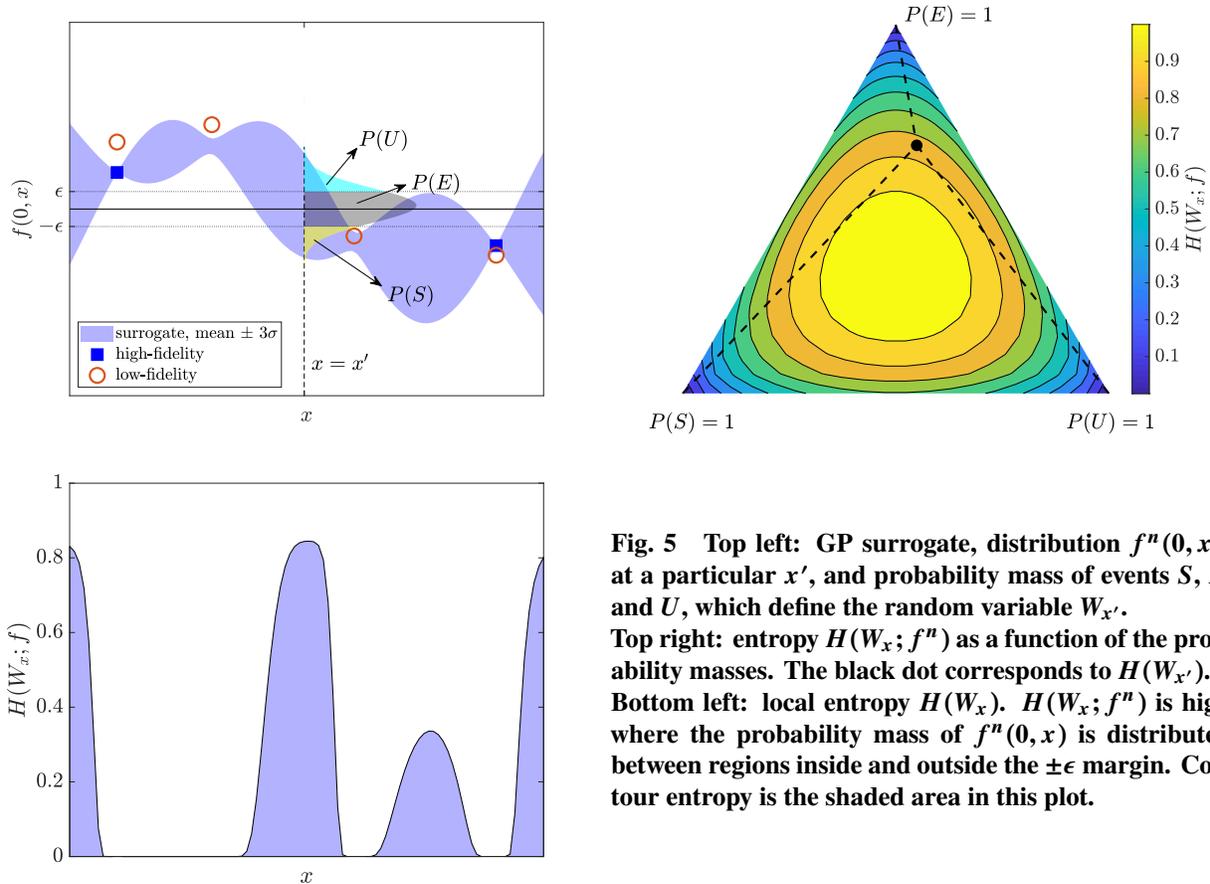
For any given condition  $\mathbf{x}$ , the aeroelastic system is stable ( $y_0(\mathbf{x}) < 0$ ), unstable ( $y_0(\mathbf{x}) > 0$ ), or in neutral equilibrium ( $y_0(\mathbf{x}) = 0$ ). The posterior surrogate model  $f^n(0, \mathbf{x})$ , conditioned on  $n$  evaluations, is a normal random variable with known mean  $m^n(0, \mathbf{x})$  and variance  $\sigma^2(0, \mathbf{x}) = \Sigma^n((0, \mathbf{x}), (0, \mathbf{x}))$ , which allows us to estimate the aeroelastic stability and quantify the uncertainty in this estimate. In order to quantify the probability of the neutral equilibrium state, we relax the definition of the flutter boundary to  $|y_0(\mathbf{x})| < \epsilon(\mathbf{x})$ , where  $\epsilon(\mathbf{x})$  is a small positive number. Then, an observation  $y(\mathbf{x})$  of  $f^n(0, \mathbf{x})$  can be classified as one of the following three events:

- $y(\mathbf{x}) < -\epsilon(\mathbf{x})$  (stable, denoted as event  $S$ ),
- $|y(\mathbf{x})| < \epsilon(\mathbf{x})$  (neutral equilibrium, denoted as event  $E$ ), or
- $y(\mathbf{x}) > \epsilon(\mathbf{x})$  (unstable, denoted as event  $U$ ).

These three events,  $S$ ,  $E$ , and  $U$ , define a discrete random variable  $W_{\mathbf{x}}$  with probability mass

$$\begin{aligned}
 P(S) &= \Phi\left(\frac{-\epsilon(\mathbf{x}) - m^n(0, \mathbf{x})}{\sigma(0, \mathbf{x})}\right), \\
 P(E) &= \Phi\left(\frac{\epsilon(\mathbf{x}) - m^n(0, \mathbf{x})}{\sigma(0, \mathbf{x})}\right) - \Phi\left(\frac{-\epsilon(\mathbf{x}) - m^n(0, \mathbf{x})}{\sigma(0, \mathbf{x})}\right), \\
 P(U) &= \Phi\left(\frac{-\epsilon(\mathbf{x}) + m^n(0, \mathbf{x})}{\sigma(0, \mathbf{x})}\right),
 \end{aligned}$$

where  $\Phi$  is the standard normal cumulative distribution function. Figure 5 illustrates events  $S$ ,  $E$ , and  $U$ , and the probability mass associated with each of them. The parameter  $\epsilon(\mathbf{x})$  influences the balance between exploration (making evaluations in regions of high uncertainty but distant from the estimated zero contour) and exploitation (making



**Fig. 5** Top left: GP surrogate, distribution  $f^n(0, x')$  at a particular  $x'$ , and probability mass of events  $S$ ,  $E$ , and  $U$ , which define the random variable  $W_{x'}$ . Top right: entropy  $H(W_x; f^n)$  as a function of the probability masses. The black dot corresponds to  $H(W_{x'})$ . Bottom left: local entropy  $H(W_x)$ .  $H(W_x; f^n)$  is high where the probability mass of  $f^n(0, x)$  is distributed between regions inside and outside the  $\pm\epsilon$  margin. Contour entropy is the shaded area in this plot.

evaluations around the estimated zero contour to refine estimate). By making  $\epsilon(x)$  proportional to  $\sigma(0, x)$ , confidence in the surrogate model is used to automatically trade-off between exploration and exploitation. If confidence is low, reflected by a large standard deviation, the margin is made wider to favor exploration. Otherwise, the margin is made narrower to favor exploitation. Reference [15] shows that  $\epsilon(x) = 2\sigma(0, x)$  offers a good compromise between exploration and exploitation, and the same value is adopted here.

The entropy of  $W_x$  measures the uncertainty in the stability of the aeroelastic system at condition  $x$ , and is given by

$$H(W_x; f^n) = -P(S) \ln P(S) - P(E) \ln P(E) - P(U) \ln P(U). \quad (8)$$

Let  $\mathcal{D}$  denote the range of parameters considered for aeroelastic flutter analysis (e.g., the flight envelope of an aerospace vehicle). To characterize the uncertainty of the flutter boundary predicted by the surrogate model we measure the *contour entropy*, defined as

$$\mathcal{H}(f^n) = \frac{1}{V(\mathcal{D})} \int_{\mathcal{D}} H(W_x; f^n) dx, \quad (9)$$

where  $V(\mathcal{D})$  denotes the volume of the parameter domain  $\mathcal{D}$ . The bottom left pane of Fig. 5 depicts the definition of

contour entropy.

### E. Selecting new evaluations

At each new iteration the multifidelity method selects which aeroelastic model  $\ell$  and aeroelastic condition  $\mathbf{x}$  to evaluate such that the uncertainty in the estimate of flutter boundary is reduced the most, per unit computational cost.

Consider the algorithm after  $n$  evaluations, with the posterior surrogate  $f^n$  and corresponding mean  $m^n$  and covariance kernel  $\Sigma^n$ . Then, the multifidelity method selects  $\ell$  and  $\mathbf{x}$  for a new evaluation by solving the following optimization problem.

$$\underset{\ell \in \{0, \dots, N_m - 1\}, \mathbf{x} \in \mathcal{D}}{\text{maximize}} \quad u(\ell, \mathbf{x}; f^n), \quad (10)$$

where

$$u(\ell, \mathbf{x}; f^n) = \frac{\mathbb{E}_y[\mathcal{H}(f^n) - \mathcal{H}(f^{n+1}) \mid \ell^{n+1} = \ell, \mathbf{x}^{n+1} = \mathbf{x}]}{p_\ell(\mathbf{x})}, \quad (11)$$

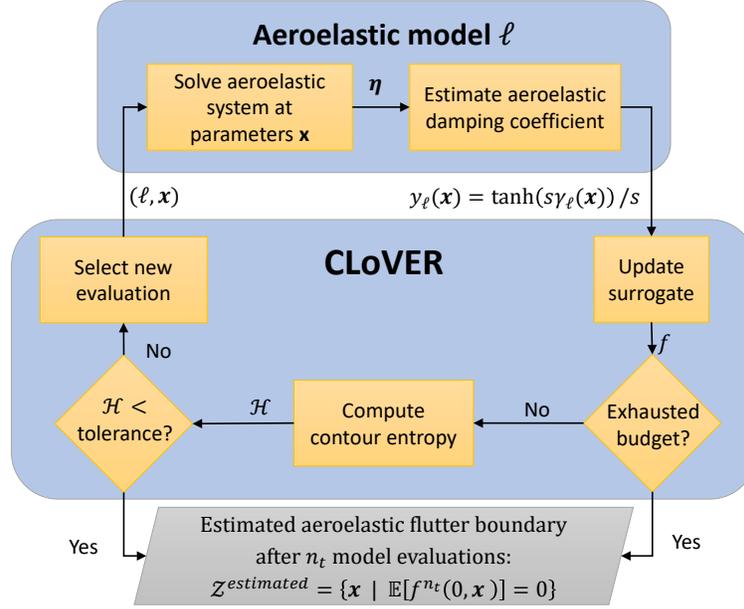
and the expectation is taken over the distribution of possible predictions,

$$y^{n+1} \sim \mathcal{N}(m^n(\ell^{n+1}, \mathbf{x}^{n+1}), \Sigma^n((\ell^{n+1}, \mathbf{x}^{n+1}), (\ell^{n+1}, \mathbf{x}^{n+1}))).$$

Note that, although contour entropy is computed as a function of the high-fidelity surrogate model only, low-fidelity observations also affect the objective function (11). Contour entropy depends on the posterior distribution of the high-fidelity GP,  $f^n(0, \mathbf{x})$ , conditioned on all available observations. In the multifidelity surrogate formulation presented in Sec. III.C, the high-fidelity GP is correlated to low-fidelity observations. Hence, as per standard rules of Gaussian process regression [36], low-fidelity observations affect the posterior distribution of the high-fidelity GP, and thus contour entropy.

There is no exact closed form expression for computing the expectation in (11), but Ref. [15] shows an approximate expression that can be computed without numerical integration. With this approximation, the utility function (11) can be evaluated as

$$u(\ell, \mathbf{x}; f^n) \approx \frac{1}{V(\mathcal{D})} \int_{\mathcal{D}} H(W_{\mathbf{x}'}; f^n) - \frac{r_\sigma(\mathbf{x}'; \ell, \mathbf{x})}{e} \sum_{i=0}^1 \sum_{j=0}^1 \exp\left(-\frac{1}{2} \left( \frac{m^n(0, \mathbf{x}') + (-1)^i \epsilon}{\hat{\sigma}(\mathbf{x}'; \ell, \mathbf{x})} + (-1)^j \beta r_\sigma(\mathbf{x}'; \ell, \mathbf{x}) \right)^2\right) d\mathbf{x}', \quad (12)$$



**Fig. 6** Schematic representation of the multifidelity method for aeroelastic flutter analysis. The algorithm CLOVER combines data from multiple aeroelastic models to train a statistical surrogate model (denoted by  $f$ ) of the aeroelastic damping coefficient. The surrogate is then used as a generative model to select which aeroelastic model and airflow condition to evaluate next.

where  $e \approx 0.577$  denotes Euler's constant,  $\beta = \Phi^{-1}(e^{-1})$ , and

$$\hat{\sigma}^2(\mathbf{x}'; \ell, \mathbf{x}) = \Sigma^{n+1}((0, \mathbf{x}'), (0, \mathbf{x}')) + \frac{(\Sigma^n((0, \mathbf{x}'), (\ell, \mathbf{x})))^2}{\Sigma^n((\ell, \mathbf{x}), (\ell, \mathbf{x}))},$$

$$r_\sigma^2(\mathbf{x}'; \ell, \mathbf{x}) = \frac{\Sigma^{n+1}((0, \mathbf{x}'), (0, \mathbf{x}'))}{\hat{\sigma}^2(\mathbf{x}'; \ell, \mathbf{x})}.$$

The integral over  $\mathcal{D}$  is computed numerically using the importance sampling approach proposed in Ref. [38].

## F. Assembling the full method

The multifidelity method for aeroelastic flutter analysis can be summarized as follows (see Fig. 6):

- 1) Compute an initial set of samples by evaluating all  $N_m$  aeroelastic models at the same values of  $\mathbf{x}$ . Use samples to compute hyperparameters and the posterior of surrogate model  $f$ .
- 2) While contour entropy is greater than set tolerance and budget is not exhausted, do:
  - 1) Determine which aeroelastic model and condition to sample next by solving the optimization problem (10).
  - 2) Evaluate the next sample at location  $\mathbf{x}^{n+1}$  using aeroelastic model  $\ell^{n+1}$ .
  - 3) Update hyperparameters and posterior of  $f$ .
- 3) Return the zero contour of  $\mathbb{E}[f^{n_t}(0, \mathbf{x})]$ , where  $n_t$  corresponds to the total number of model evaluations.

## IV. Aeroelastic flutter analysis of Isogai case A

### A. Problem description

We apply the multifidelity method to locate the flutter boundary of the aeroelastic system “case A” introduced by Isogai [25, 39, 40]. This system is an instance of a typical section model, which represents the aeroelastic behavior of a rigid airfoil supported by linear translation and torsion springs [41, 42]. The motion of the typical section is represented by two degrees of freedom: pitch angle ( $\theta$ ) and vertical translation ( $\xi$ ) – see Fig. 7 for conventions. The aeroelastic dynamics of this model are governed by (in dimensionless form):

$$M\ddot{\boldsymbol{\eta}} + K\boldsymbol{\eta} = \frac{V^2}{\pi} Q_a(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}, \ddot{\boldsymbol{\eta}}, \boldsymbol{\eta}_a, \dot{\boldsymbol{\eta}}_a), \quad (13a)$$

$$\dot{\boldsymbol{\eta}}_a = \mathcal{A}(M_\infty, \boldsymbol{\eta}_a, \boldsymbol{\eta}, \dot{\boldsymbol{\eta}}), \quad (13b)$$

where

$$\boldsymbol{\eta} = \begin{Bmatrix} \xi \\ \theta \end{Bmatrix}, \quad M = \begin{bmatrix} 1 & x_\theta \\ x_\theta & r_\theta^2 \end{bmatrix}, \quad K = \begin{bmatrix} (\omega_h/\omega_\theta)^2 & 0 \\ 0 & r_\theta^2 \end{bmatrix}, \quad Q_a = \begin{Bmatrix} -c_\ell \\ 2c_m \end{Bmatrix}. \quad (14)$$

In the equation above  $x_\theta = x_{cg} - x_{ea}$  is the static imbalance, where  $x_{cg}$  denotes the position of the center of gravity and  $x_{ea}$  the position of the elastic axis, both measured from the mid-chord and non-dimensionalized with respect to the semi-chord  $b$ . Furthermore,  $r_\theta$  is the radius of gyration about the elastic axis,  $\omega_h$  and  $\omega_\theta$  are the uncoupled natural angular frequencies of vibration of the heave and pitch modes,  $c_\ell$  denotes the lift coefficient, and  $c_m$  denotes the pitching moment coefficient with respect to the elastic axis. Flutter speed is represented in nondimensional form by the speed index  $V_\mu = V_\infty/(\sqrt{\mu}b\omega_\theta)$ , where  $\mu = w/(\pi\rho_\infty b^2)$  is the mass ratio,  $\rho_\infty$  and  $V_\infty$  are the density and speed at freestream, respectively, and  $w$  is the mass of the airfoil per unit span. The Isogai case A problem [39] is based on a NACA64A010 airfoil and the following set of parameters:

$$\mu = 60, \quad r_\theta^2 = 3.48, \quad \omega_h = 100 \text{ rad/s}, \quad \omega_\theta = 100 \text{ rad/s}, \quad x_{cg} = -0.2, \quad x_{ea} = -2.$$

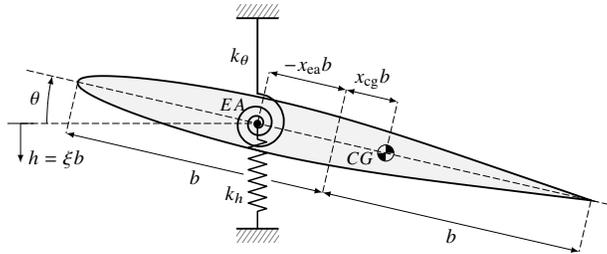


Fig. 7 Typical section model.

We assume the airflow is in the high Reynolds number regime, such that it is reasonable to neglect viscous effects. Two models of the airflow dynamics are considered—Euler equations and linearized potential flow—resulting in aeroelastic models of varying fidelity, which are described below. Furthermore, the aeroelastic flutter boundary is described as a function of two aerodynamic parameters: Mach number ( $M_\infty$ ) and speed index ( $V_\mu$ ). The parameter space is set to  $(M_\infty, V_\mu) \in [0.6, 0.9] \times [0.4, 2.0]$ , which is known to include the transonic dip phenomenon for this particular problem.

## B. Aeroelastic models

### 1. Overview

We use four aeroelastic models to locate the flutter boundary of the Isogai case A problem. Three of these models are based on integrating (13) in time using the open source software SU2 [43], employing the Euler equations as the formulation of the airflow dynamics. Further details are in Sec. IV.B.2. The fourth model solves (13) in the frequency domain using a doublet lattice model discretization of the linearized potential flow formulation. This model is discussed in Sec. IV.B.3. The list below summarizes the aeroelastic models used in the present multifidelity analysis<sup>‡</sup>:

- High-fidelity model (HFM): time integration using SU2, Euler equations, fine mesh (8606 points). Integration time:  $20\pi/\omega_\theta$ . Computational cost:  $\approx 11,870$  CPU-s.
- Low-fidelity model 1 (LFM1): time integration using SU2, Euler equations, fine mesh (8606 points). Integration time:  $6\pi/\omega_\theta$ . Computational cost:  $\approx 3,950$  CPU-s.
- Low-fidelity model 2 (LFM2): time integration using SU2, Euler equations, coarse mesh (4572 points). Integration time:  $6\pi/\omega_\theta$ . Computational cost:  $\approx 1,560$  CPU-s.
- Low-fidelity model 3 (LFM3): frequency domain using p-k method, doublet lattice method (30 panels). Computational cost:  $\approx 8$  CPU-s.

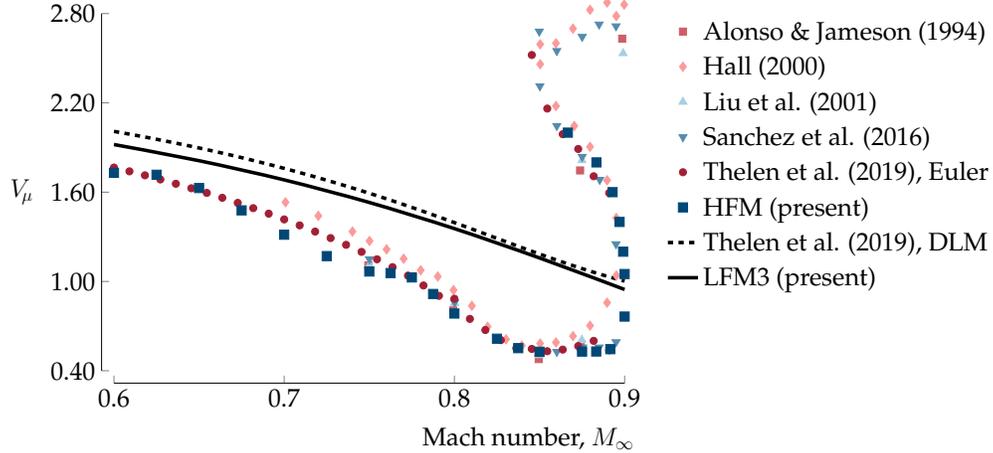
### 2. Aeroelastic models based on the Euler equations (HFM, LFM1, and LFM2)

The aeroelastic models HFM, LFM1, and LFM2 use the Euler equations to model the airflow dynamics. The Euler equations are a set of conservation laws that represent the dynamics of inviscid, compressible, and rotational flows, and include non-linear effects such as shock waves. The open source software SU2 [43] is used to advance the structural and aerodynamic states in time by solving (13) with an implicit Euler discretization. The time step is set to  $\Delta t = 2\pi/100\omega_\theta$  (1/100th of the period of the undamped pitch mode of vibration), with 20 sub-iterations per time step. SU2 uses a finite-volume discretization of the conservation laws, and the second order Jameson-Schmidt-Turkel (JST) flux scheme with the Venkatakrisnan limiter is used in the present investigation.

Evaluations of these models start by solving the steady Euler equations for a fixed angle-of-attack of one degree.

---

<sup>‡</sup>Computational cost measured on a Quad-Core Intel® Xeon™ processor E5-1620, 3.60 GHz, 10 MB Cache, 32 GB RAM. Simulations are carried out on a single core.



**Fig. 8 Comparison of flutter boundary estimated by present HFM and LFM3 models to results reported in the literature. Thelen et al. [18] report results based on the Euler equations and linearized potential flow (DLM). All other literature results are based on the Euler equations.**

The steady aerodynamic solution is then used as initial condition to the unsteady aeroelastic problem. The matrix pencil method [44]<sup>§</sup> is used to estimate damping and frequency from the time history of pitch angle and vertical translation. Each of these degrees of freedom is used independently to estimate damping and frequency, and the value corresponding to the largest damping determines the aeroelastic damping coefficient. Following the suggestion of Jacobson et al. [44] we discard an initial portion of the time history to eliminate transient effects due to the initial condition. In HFM the initial 300 time steps (30%) are discarded, whereas in LFM1 and LFM2 the initial 50 time steps (16%) are discarded.

A mesh convergence study was conducted by evaluating the flutter speed index at several Mach numbers with three increasingly refined meshes. The flutter speed index in this study was determined using the bisection method. Table 1 shows the convergence of flutter speed index for  $M_\infty = 0.60, 0.75, \text{ and } 0.90$ . The variation of flutter speed index from the medium to the fine mesh is lower than 2% in all cases, which we consider an acceptable threshold for convergence.

Furthermore, we apply the bisection method to the HFM (fine mesh) to locate 24 discrete points along the flutter boundary. These points are listed in the Appendix. Figure 8 compares these predictions to results reported in the literature [1, 18, 40, 45, 46]. The results computed with the HFM are within the range of results computed with other solvers based on the Euler equations.

The parameter space selected in the present analysis,  $(M_\infty, V_\mu) \in [0.6, 0.9] \times [0.4, 2.0]$ , includes the lower part of the flutter boundary shown in Fig. 8, which is the most relevant for practical purposes. Although investigations were not conducted beyond this range of interest, the multifidelity method is expected to work well as long as aeroelastic methods can be used to produce reliable estimates of aeroelastic damping coefficient.

<sup>§</sup>Jacobson et al. [44] compared several techniques to estimate damping and frequency from time history of oscillations of aeroelastic systems and found the matrix pencil method to be the most robust.

**Table 1** Mesh convergence study. The variation of flutter speed index between medium and fine meshes is lower than 2%.

mesh	# nodes	$V_{\mu}^{flutter}$ @ $M_{\infty} = 0.60$	$V_{\mu}^{flutter}$ @ $M_{\infty} = 0.75$	$V_{\mu}^{flutter}$ @ $M_{\infty} = 0.90$
coarse	4572	1.738	1.090	0.6922
medium	6738	1.735	1.088	0.7500
fine	8606	1.730	1.067	0.7641

### 3. Aeroelastic model based on linearized potential flow (LFM3)

The aeroelastic model LFM3 is based on the linearized potential flow formulation of the airflow dynamics. This formulation assumes that the airflow can be described by inviscid, isentropic, and irrotational small disturbances about a uniform velocity field. This formulation accounts for compressibility effects, but neglects effects such as airfoil shape and shock waves and therefore cannot capture the transonic dip. The aerodynamic forces are computed by assuming harmonic oscillations of pitch angle and vertical displacement, and solving the Possio’s integral equation [41] for the corresponding distribution of acceleration potential doublets. In our implementation, Possio’s integral equation is solved using a collocation method analogous to the doublet-lattice method (DLM) [12] used in three-dimensional problems. LFM3 discretizes the doublet distribution using 30 panels along the length of the airfoil. The software used to compute aerodynamic forces is available at <https://github.com/anmarques/DLM2D>.

The aeroelastic damping coefficient is computed by solving (13) in the frequency domain using the  $p$ - $\kappa$  method [47]. Figure 8 compares the flutter boundary estimated by LFM3 and the one computed by Thelen et al. [18] using the DLM implementation of the software ASTROS [48]. The discrepancy in flutter speed index is below 5%, which we consider satisfactory for the purposes of the present investigation. Furthermore, we conjecture that the discrepancy is mostly due to the fact that the DLM formulation of ASTROS is valid for three-dimensional wings, and the results of Thelen et al. [18] are computed with a large aspect ratio wing instead of a truly two-dimensional airfoil.

### C. Initialization, surrogate models, and hyperparameters

We initialize the surrogate models using three data points<sup>¶</sup> located along the flutter boundary predicted by LFM3. These points are estimated using the bisection method, resulting in the initial design set displayed in Table 2. Estimating these points with the inexpensive LFM3 takes approximately 30s. All four aeroelastic models are evaluated at the initial design set, and the resulting data is used to initialize the multifidelity surrogate model. The upper left frame of Fig. 10 shows the surrogate model trained at the initial design set.

The prior knowledge about the flutter boundary is expressed using a linear mean function fitted to the initial data for

<sup>¶</sup>Other initialization strategies can also be adopted. For instance, the initial data may be selected to reflect the agreement between aeroelastic models.

**Table 2 Initial design set. Three points located along the flutter boundary predicted by LFM3.**

$M_\infty$	$V_\mu$
0.60	1.9200
0.75	1.5309
0.90	0.9460

the HFM:

$$\mu_0(\mathbf{x}) = -0.38 + 0.30x_1 + 0.12x_2,$$

where  $\mathbf{x} = \{M_\infty, V_\mu\}$ . This function captures the general trend that the aeroelastic damping coefficient increases with speed and Mach number, and that the flutter speed decreases as Mach number increases. This trend is not valid in regions of significant non-linear behavior, but the method proposed here is able to correct the prior model as new data is observed. We further assume a zero mean function as prior for the discrepancies between the HFM and low-fidelity models.

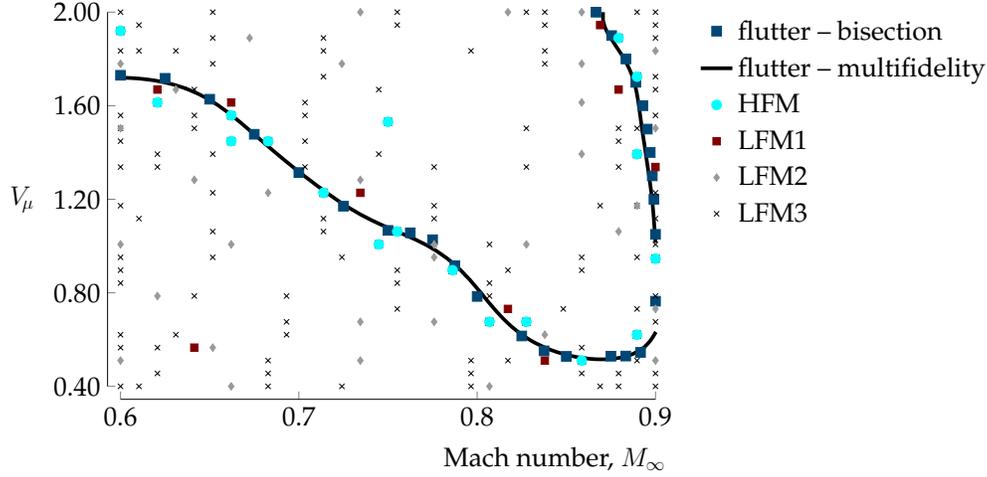
We use covariance kernels of the squared exponential type [36],

$$\Sigma_\ell(\mathbf{x}_p, \mathbf{x}_q) = \sigma_\ell^2 \exp\left(-0.5(\mathbf{x}_p - \mathbf{x}_q)^T M_\ell^2 (\mathbf{x}_p - \mathbf{x}_q)\right), \quad (15)$$

where  $M_\ell$  is a diagonal matrix whose entries correspond to the inverse of correlation lengths. The initial values of hyperparameters for the covariance kernels are listed in Table 3. These values are updated using a maximum likelihood estimate whenever the method makes a new HFM evaluation. However, to avoid spurious hyperparameter estimates due to insufficient data (especially at the start of the iterations), we limit the hyperparameters to lie within 50% of the initial values.

**Table 3 Initial set of hyperparameters for covariance kernels.**

model	variance ( $\sigma_\ell^2$ )	correlation lengths ( $1/\text{diag}(M_\ell)$ )
HFM ( $\ell = 0$ )	$5 \times 10^{-3}$	{0.050, 0.200}
LFM1 ( $\ell = 1$ )	$7 \times 10^{-5}$	{0.025, 0.143}
LFM2 ( $\ell = 2$ )	$7 \times 10^{-5}$	{0.025, 0.143}
LFM3 ( $\ell = 3$ )	$5 \times 10^{-4}$	{0.050, 0.143}



**Fig. 9** Model evaluations used by the multifidelity method (selected from a  $30 \times 30$  grid in parameter space), and resulting flutter boundary estimate. The estimated flutter boundary is in good agreement with points computed with the bisection method.

#### D. Optimization method

As discussed in Sec. III, the multifidelity method selects new evaluations by solving the optimization problem (10) at every iteration. Here this optimization problem is solved by performing a search on a uniform  $4 \times 30 \times 30$  grid in the space  $(\ell, M_\infty, V_\mu) \in \{0, 1, 2, 3\} \times [0.6, 0.9] \times [0.4, 2]$ .

#### E. Flutter boundary prediction

The multifidelity method proposed here achieved the stopping criterion  $\mathcal{H} = 0.01$  after 172 iterations, at a cost<sup>¶</sup> equivalent to 20 HFM evaluations. Figure 9 shows that the flutter boundary estimated by the multifidelity method is in good agreement with points computed using the bisection method (available in the Appendix). Out of the 18 bisection points located in the lower portion of the flutter boundary, the error produced by the multifidelity method is below 5% at 17 points, and around 17% at  $M_\infty = 0.9$ . The larger error at  $M_\infty = 0.9$  occurs due to an abrupt change in flutter speed that requires additional model evaluations to be captured. Although not shown here, the error at  $M_\infty = 0.9$  becomes smaller than 5% if the multifidelity method is allowed to advance to 176 iterations, with a cost of 22 HFM evaluations. In addition, the cost of the multifidelity method is 85% smaller than the cost of computing the 18 reference points with the bisection method, which required 129 HFM evaluations.

Figure 9 also shows where each model is evaluated, demonstrating how the multifidelity method allocates computational resources. The cheapest model, LFM3, is used to explore most of the parameter space and is evaluated 128 times. The other models are considerably more expensive than LFM3, and the multifidelity method uses them more

<sup>¶</sup>The cost listed above accounts only for evaluations of aeroelastic models. The cost of using CLOVER to select which aeroelastic model and data point to evaluate is approximately 20 CPU-s in the present investigation. This cost is of the same order as the cost of evaluating the cheapest aeroelastic model, LFM3, but negligible with respect to the cost of evaluating higher-fidelity models that dominate the overall cost of aeroelastic analysis.

sparingly. LFM2 is evaluated 52 times, with evaluations used both to learn the location of the flutter boundary and to gain confidence that flutter does not occur in other regions of the parameter space. LFM1 is evaluated 28 times, with 26 evaluations located in the vicinity of the flutter boundary.\*\* Finally, all 18 evaluations of HFM are very close to the flutter boundary, which allows the method to make accurate predictions.

Figure 10 shows the distribution of the local entropy  $H(W_{\mathbf{x}})$  at several snapshots along the evolution of the iterative process. The initial setup reflects our choice of prior model with a linear mean function. As expected, entropy is high everywhere in the parameter space, with exception of regions surrounding the three points used to initialize the calculations. The multifidelity method initially uses the cheapest models LFM2 and LFM3 to explore the parameter space, reducing entropy in most locations and narrowing down the regions where the flutter boundary is likely to occur. Then, the method balances further exploration (mostly with LFM2 and LMF3) with exploitation of the flutter boundary. Note that entropy is not guaranteed to decrease at every iteration. For instance, at iteration 77 entropy is large in regions where it was low at iteration 57. This behavior reflects changes in the hyperparameters (updated every time an HFM evaluation is made), and new observations that indicate the presence of flutter at the upper right corner of the parameter space.

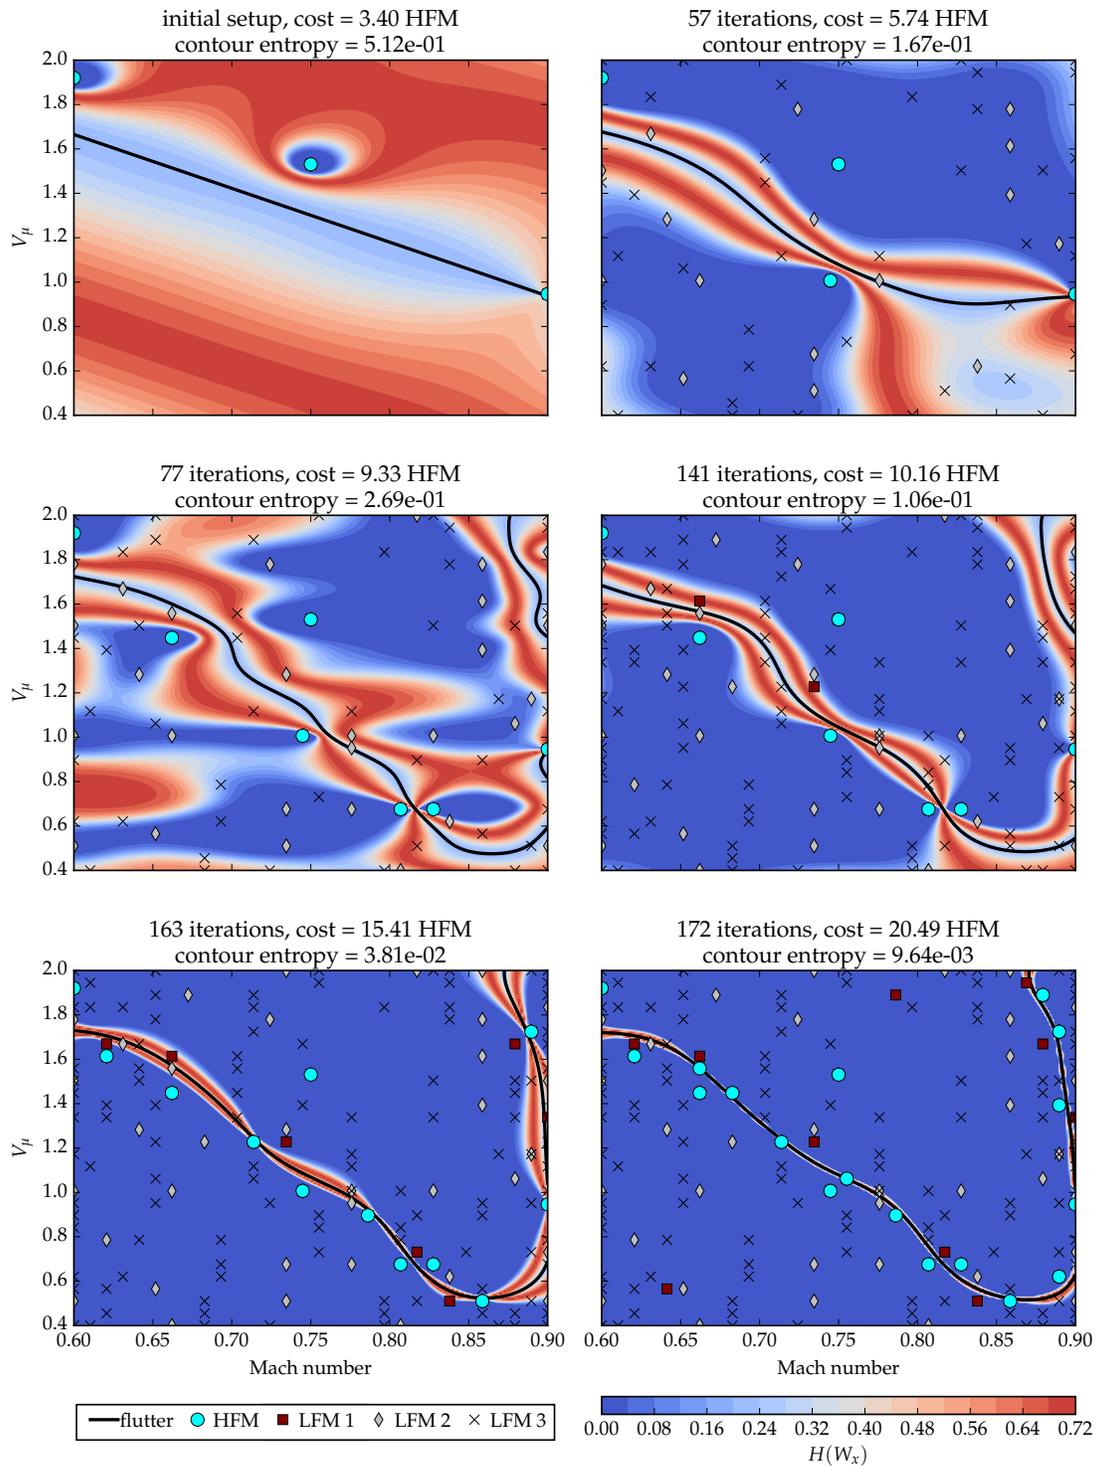
We can also observe from Fig. 10 that the multifidelity method produces reasonable estimates of the flutter boundary in as little as 77 iterations, at a cost of 9 HFM evaluations. We quantify the quality of the flutter boundary estimate by averaging the absolute error over the 18 bisection points in the lower portion of the flutter boundary,

$$\frac{1}{18} \sum_{i=1}^{18} |V_{\mu}^{bisection}(\mathbf{x}_i) - V_{\mu}^{multifidelity}(\mathbf{x}_i)|. \quad (16)$$

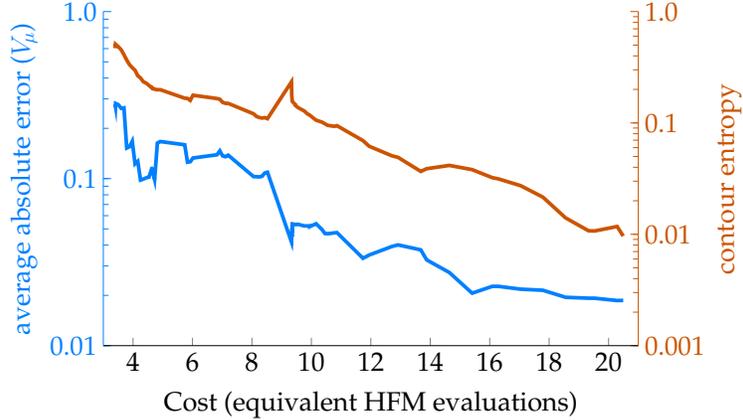
Figure 11 shows the evolution of the average absolute error as a function of the computational cost, along with the contour entropy. The average absolute error does not vary abruptly as contour entropy approaches 0.01, indicating that the estimate produced with the stopping criterion  $\mathcal{H} = 0.01$  is robust for this particular problem. However, reasonable estimates are possible at lower cost.

---

\*\*LFM1 uses the same solver and grid as HFM. The difference between the models is that LFM1 uses a shorter integration time to estimate damping. Hence, whenever both models are evaluated at the same location, it is possible to obtain HFM and LFM1 estimates using a single CFD simulation. This is reflected in the present results.



**Fig. 10** Entropy distribution at several snapshots along the iterative process. Model evaluations are selected such that entropy is reduced the most, in expectation, per unit cost.



**Fig. 11 Evolution of the absolute error in the estimate of the flutter boundary and contour entropy. The absolute error is averaged over the 18 bisection points on the lower portion of the flutter boundary.**

## V. Discussion

In this section we discuss in detail some topics that can influence the application of our multifidelity method in general aeroelastic problems. In Sec. V.A, we discuss about the choice of stopping criterion, and the effect of noisy aeroelastic models. In Sec. V.B, we discuss how to address failed evaluations of aeroelastic models and their potential consequences. In Sec. V.C, we address the topic of selecting aeroelastic models, including the possibility of models based on time-domain and frequency-domain formulations. Finally, in Sec. V.D, we discuss the challenges of scaling the multifidelity method to a large number of parameters.

### A. Stopping criterion and effect of noise

The multifidelity method is driven by reducing contour entropy. Hence, it is natural to define a stopping criterion based on contour entropy itself. However, contour entropy is influenced by the choice of statistical surrogate models (e.g., covariance kernels and hyperparameters) and thus is not a universal indicator of performance. If different statistical surrogate models are used, investigating the relationship between contour entropy and other performance indicators—such as error (if available), changes in predicted flutter boundary, and changes in hyperparameters—is advisable.

In addition, using contour entropy as a stopping criterion guarantees that the multifidelity method converges even when a flutter boundary is not present in the parameter space. In this scenario, the method will explore the parameter space using a combination of aeroelastic models of varying fidelities, reducing the overall contour entropy as it gains confidence that flutter does not occur in any region of the parameter space until the stopping criterion is satisfied.

Noisy model evaluations can also impact the choice of stopping criterion. Although not discussed in Sec. III, the multifidelity surrogate model can incorporate random Gaussian noise if an estimate of noise variance is available. This subject is discussed in Ref. [15]. In general, the effect of noise is to reduce the overall confidence in model evaluations,

requiring more data to make accurate predictions. Furthermore, if the HFM is noisy, then the overall contour entropy is limited by a minimum level resulting from noise. The stopping criterion must be selected above this level for the analysis to converge.

## **B. Failed model evaluations**

Numerical solvers can fail for several reasons (e.g., segmentation fault, overflow errors), leading to model evaluations that produce unreliable results or no results at all. If a model evaluation fails such that no meaningful estimates can be made (e.g., early segmentation fault), the evaluation should be ignored, and the data point removed from the set of possible solutions for the evaluation selection problem. In principle, we expect the multifidelity method not to be significantly affected if the instances of failed evaluations are few relative to the total number of evaluations.

The only failure mode we observed was divergent oscillations that achieve nonphysical values of pitch angle, sometimes interrupting execution early due to overflow errors. Such cases occur when the system is strongly divergent, and hence where accurate estimates of damping are not needed. This situation was handled by cropping the data used for damping estimates if pitch angle became greater than 5 degrees in magnitude, and letting the definition of quantity of interest introduced in Sec. III.B naturally saturate to an upper limit.

## **C. Selecting aeroelastic models**

The physical and mathematical formulations used to describe the aeroelastic dynamics have significant effects on the fidelity and the cost of aeroelastic models, and the parametrization of the aeroelastic problem. Below we discuss how such factors affect the selection of aeroelastic models for multifidelity analysis of flutter.

One important distinction between aeroelastic models is whether they are based on time-domain or frequency-domain formulations. The multifidelity method presented here can incorporate information from aeroelastic models based on both formulations. However, because time-domain and frequency-domain models depend on different sets of physical variables, they require distinct parametrizations of the aeroelastic problem. In the example of Sec. IV, the lowest-fidelity model (LFM3) is based on a frequency-domain formulation, whereas the other three models (LFM1, LFM2, and HFM) are based on a time-domain formulation. Because the models based on a time-domain formulation are more expensive to evaluate, the flutter problem was parametrized in a way to favor the time-domain formulation. Estimating aeroelastic damping for a given combination of Mach number ( $M_\infty$ ) and flutter speed index ( $V_\mu$ ) requires a single evaluation of the numerical solver associated with the time-domain formulation, whereas multiple evaluations of the frequency-domain DLM solver are needed to estimate damping for LFM3. This approach is feasible because the DLM solver is cheap to evaluate.

Adopting a high-fidelity model based on a frequency-domain formulation (e.g., aeroelastic models discussed in [17, 49]) is also possible, but it requires a different parametrization of the flutter problem to reduce the number of

evaluations of the frequency-domain aerodynamic solver. For instance, the parameter space can be defined in terms of Mach number ( $M_\infty$ ) and reduced frequency ( $\kappa = \omega b/V_\infty$ ), and the quantity of interest be the damping estimate produced by the  $\kappa$  method [47]<sup>††</sup>. With this parametrization, evaluating a data point requires one evaluation of the frequency-domain aerodynamic solver per structural mode involved in the analysis. Furthermore, the resulting flutter boundary can still be expressed in terms of the flutter speed index since this quantity is also computed when evaluating the frequency-domain aeroelastic solver.

Another important consideration for model selection is the availability of models of intermediate fidelity. Our experience shows that, when the cost of the high-fidelity model is significantly higher than the lowest-fidelity model, including models of intermediate fidelity and cost is beneficial to the multifidelity analysis. In Sec. IV, the multifidelity method uses the lowest-fidelity model (LFM3) to explore most of the parameter space, but it requires higher-fidelity evaluations to learn in each regions of the parameter space LFM3 can be trusted. Several evaluations of the intermediate-fidelity model LFM2 are used for this purpose. If LFM1 and LFM2 were not part of the model spectrum, the multifidelity method would require several additional evaluations of HFM to gain confidence in the predictions of LFM3, which would have a significant impact on the overall cost.

The spectrum of models used in Sec. IV reflects the experience of the authors with aeroelastic simulations to achieve a reasonable trade-off between fidelity and cost. We include an aeroelastic model based on simplified physics (LFM3) because it has very low computational cost and offers relatively accurate predictions in regions of the parameter space where the aeroelastic system is not dominated by non-linear effects. At the other end of the spectrum, we select a high-fidelity model (HFM) based on a refined time integration of the non-linear aeroelastic dynamics to guarantee the accuracy of the results. Given the large cost discrepancy between LFM3 and HFM (1:1480), we include intermediate models by reducing the spatial resolution and integration time used in HFM. The coarse-grid model (LFM2) achieves a significant cost reduction with respect to HFM (1:8). We did not consider coarsening the grid further because LFM2 predicts a more pronounced transonic dip phenomenon than HFM, and we conjecture that further cost reduction might not justify losses in accuracy. Furthermore, the integration time of LFM1 and LFM2 corresponds to approximately three periods of oscillation of the aeroelastic system. We did not consider reducing this integration time further because it is only slightly higher than the minimum of two cycles needed to determine if the system is stable or unstable. Finally, LFM1 is simply a time-truncated version of HFM. Hence, every evaluation of HFM offers a free “evaluation” of LFM1. For this reason, we find that including LFM1 into the spectrum of models is beneficial, although this model is queried less often by the multifidelity method.

---

<sup>††</sup>The damping computed by the  $\kappa$  method is not equivalent to the damping computed by the  $p$ - $\kappa$  method. However, both methods estimate the same flutter boundary. See Ref. [47] for further details.

#### D. Aeroelastic problems with large number of parameters

Aeroelastic systems can be influenced by a large number of parameters, resulting in aeroelastic analysis in high-dimensional spaces. The main challenge of high-dimensional problems is that the number of samples needed to cover a region of space grows exponentially with the number of dimensions (*curse of dimensionality*). The multifidelity method presented here mitigates this issue by sampling expensive high-fidelity models mostly along a lower dimensional hypersurface corresponding to the flutter boundary. In addition, physical problems often present a lot of structure and smoothness in parameter space, which can be exploited in the definition of Gaussian process priors to reduce the overall number of samples even further.

Another challenge associated with increasing the number of samples is the cost of training a Gaussian process surrogate of the quantity of interest. The cost of training GPs scales as  $n^3$ , where  $n$  is the number of samples. For problems involving large number of samples ( $n \gg 10,000$ ), GP sparsification techniques [36, 50] can be used to trade some accuracy for significant savings in computational time.

In summary, one should carefully choose the parameters to include in the flutter analysis to maintain the cost within reasonable limits. In principle, sophisticated modeling and numerical techniques (GP priors and sparsification) can be used to mitigate the effects of the curse of dimensionality in the context of the multifidelity method, but high-dimensional problems are inherently expensive. We have not explored such techniques in our analysis.

## VI. Conclusion

We presented a multifidelity method that locates aeroelastic flutter boundaries by combining information from low- and high-fidelity aeroelastic models. We demonstrated through an example of the typical section model that this method can result in significant computational savings with respect to a naive exploration of the space of aeroelastic configurations (e.g., bisection method). In our example, the reduction in cost was about 85% with an average absolute error in flutter speed index of the order of 0.01, when compared to using the bisection method with the high-fidelity model only.

The proposed method uses an active learning mechanism to select evaluations of different aeroelastic models, optimally allocating computational resources in order to produce an accurate estimate of flutter. As shown in the typical section example, expensive models are only evaluated in the vicinity of the actual flutter boundary, resulting in accurate estimates with relatively few expensive evaluations.

Furthermore, the multifidelity method can combine information from aeroelastic models based on both time-domain and frequency-domain formulations. In the example presented here, we combine information from a low-fidelity model based on the frequency-domain doublet lattice method with three aeroelastic models based on integrating coupled unsteady structural dynamics and Euler equations in time.

We see a few possibilities for future work. One possibility is allowing the method to select new model evaluations in

batches, taking full advantage of parallel computations. Similar ideas have been explored in the context of Bayesian optimization and may be beneficial for aeroelastic analysis. Another possibility is investigating Gaussian process priors tailored for aeroelastic applications. As generally occurs with algorithms based on GP surrogates, the performance of the multifidelity method can be significantly affected by the choice of priors. Creating GP priors tailored for aeroelastic applications may allow the method to use information from multiple aeroelastic models more efficiently.

## Appendix

Table A.1 lists the points used as reference solution in the estimate of the flutter boundary of the Isogai case A problem. These points were estimated using the bisection method with the high-fidelity model described in Sec. IV.B.2.

**Table A.1 Points along the flutter boundary of the Isogai case A problem.**

Lower portion of flutter boundary		Upper portion of flutter boundary	
Mach number	$V_\mu$	Mach number	$V_\mu$
0.6000	1.7301	0.9000	1.0499
0.6250	1.7171	0.8991	1.2000
0.6500	1.6277	0.8982	1.3000
0.6750	1.4779	0.8971	1.4000
0.7000	1.3142	0.8954	1.5000
0.7250	1.1704	0.8929	1.6000
0.7500	1.0668	0.8890	1.7000
0.7625	1.0565	0.8834	1.8000
0.7750	1.0270	0.8754	1.9000
0.7875	0.9159	0.8666	2.0000
0.8000	0.7845		
0.8250	0.6154		
0.8375	0.5519		
0.8500	0.5281		
0.8750	0.5289		
0.8833	0.5296		
0.8916	0.5451		
0.9000	0.7641		

## Funding Sources

This work was supported in part by the U.S. Air Force Center of Excellence on Multi-Fidelity Modeling of Rocket Combustor Dynamics, Award FA9550-17-1-0195, and by the AFOSR MURI on Managing Multiple Information Sources of Multi-Physics Systems, Awards FA9550-15-1-0038 and FA9550-18-1-0023.

## Acknowledgments

We are thankful to Andrew Thelen and Prof. Leifur Leifsson for sharing their results for the Isogai case A problem. We are also thankful to Prof. Graeme Kennedy for allowing us to use his implementation of the matrix pencil method.

## References

- [1] Liu, F., Cai, J., Y. Z., Tsai, H. M., and Wong, A. S. F., “Calculation of Wing Flutter by a Coupled Fluid-Structure Method,” *Journal of Aircraft*, Vol. 38, No. 2, 2001, pp. 334–342. doi:10.2514/2.2766.
- [2] Shang, J. S., “Three decades of accomplishments in computational fluid dynamics,” *Progress in Aerospace Sciences*, Vol. 40, No. 3, 2004, pp. 173 – 197. doi:10.1016/j.paerosci.2004.04.001.
- [3] Farhat, C., “CFD on moving grids: from theory to realistic flutter, maneuvering, and multidisciplinary optimization,” *International Journal of Computational Fluid Dynamics*, Vol. 19, No. 8, 2005, pp. 595–603. doi:10.1080/10618560500510579.
- [4] Datta, A. D., Sitaraman, J., Chopra, I., and Baeder, J. D., “Consistent rational-function approximation for unsteady aerodynamics,” *Journal of Aircraft*, Vol. 28, No. 9, 1991, pp. 545–552. doi:10.2514/3.46062.
- [5] Bazilevs, Y., Calo, V. M., Hughes, T. J. R., and Zhang, Y., “Isogeometric fluid-structure interaction: theory, algorithms, and computations,” *Computational Mechanics*, Vol. 43, No. 1, 2008, pp. 3–37. doi:10.1007/s00466-008-0315-x.
- [6] Hou, G., Wang, J., and Layton, A., “Numerical Methods for Fluid-Structure Interaction — A Review,” *Communications in Computational Physics*, Vol. 12, No. 2, 2012, pp. 337—377. doi:10.4208/cicp.291210.290411s.
- [7] Smith, M. J., Lim, J. W., van der Wall, B. G., Baeder, J. D., Biedron, R. T., Boyd, D. D., Jayaraman, B., Jung, S. N., and Min, B.-Y., “The HART II international workshop: an assessment of the state of the art in CFD/CSD prediction,” *CEAS Aeronautical Journal*, Vol. 4, No. 4, 2013, pp. 345–372. doi:10.1007/s13272-013-0078-8.
- [8] Heeg, J., “Overview and Lessons Learned from the Aeroelastic Prediction Workshop,” *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA, Boston, MA, 2013. doi:10.2514/6.2013-1798.
- [9] Heeg, J., Chwalowski, P., Raveh, D. E., Jirasek, A., and Dalenbring, M., “Overview and Data Comparisons from the 2<sup>nd</sup> Aeroelastic Prediction Workshop,” *34th AIAA Applied Aerodynamics Conference (AIAA Aviation Forum)*, AIAA, Washington, DC, 2016. doi:10.2514/6.2016-3121.
- [10] Farhat, C., *CFD-Based Nonlinear Computational Aeroelasticity*, American Cancer Society, 2017, pp. 1–21. doi:10.1002/9781119176817.ecm2063.

- [11] Miller, B. A., and McNamara, J. J., “Efficient Fluid-Thermal-Structural Time Marching with Computational Fluid Dynamics,” *Journal of Aircraft*, Vol. 56, No. 9, 2018, pp. 3610–3621. doi:10.2514/1.J056572.
- [12] Albano, E., and Rodden, W. P., “A Doublet-Lattice Method for Calculating Lift Distributions on Oscillating Surfaces in Subsonic Flows,” *AIAA Journal*, Vol. 7, No. 2, 1969, pp. 279–285. doi:10.2514/3.5086.
- [13] Rodden, W. P., Giesing, J. P., and Kálmán, T. P., “Refinement of the nonplanar aspects of the subsonic doublet-lattice lifting surface method,” *Journal of Aircraft*, Vol. 9, No. 1, 1972, pp. 69–73. doi:10.2514/3.44322.
- [14] Giesing, J. P., Kálmán, T. P., and Rodden, W. P., “Subsonic Steady and Oscillatory Aerodynamics for Multiple Interfering Wings and Bodies,” *Journal of Aircraft*, Vol. 9, No. 10, 1972, pp. 693–702. doi:10.2514/3.59066.
- [15] Marques, A., Lam, R., and Willcox, K., “Contour location via entropy reduction leveraging multiple information sources,” *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Curran Associates, Inc., 2018, pp. 5222–5232.
- [16] Dribusch, C., Missoum, S., and Beran, P., “A multifidelity approach for the construction of explicit decision boundaries: application to aeroelasticity,” *Structural and Multidisciplinary Optimization*, Vol. 42, No. 5, 2010, pp. 693–705. doi:10.1007/s00158-010-0516-8.
- [17] Timme, S., and Badcock, K. J., “Transonic Aeroelastic Instability Searches Using Sampling and Aerodynamic Model Hierarchy,” *AIAA Journal*, Vol. 49, No. 6, 2011, pp. 1191–1201. doi:10.2514/1.j050509.
- [18] Thelen, A. S., Leifsson, L. T., and Beran, P. S., “Aeroelastic Flutter Prediction using Multi-fidelity Modeling of the Aerodynamic Influence Coefficients,” *AIAA Scitech 2019 Forum*, AIAA, 2019. doi:10.2514/6.2019-0609.
- [19] Roger, K. L., “Airplane Math Modeling Methods for Active Control Design,” Tech. Rep. AGARD–CP–228, Aug 1977.
- [20] Abel, I., “An Analytical Technique for Predicting the Characteristics of a Flexible Wing Equipped with an Active Flutter-Suppression System and Comparison with Wind-Tunnel Data,” Tech. Rep. NASA TP–1367, Feb 1979.
- [21] Eversman, W., and Tewari, A., “Consistent rational-function approximation for unsteady aerodynamics,” *Journal of Aircraft*, Vol. 28, No. 9, 1991, pp. 545–552. doi:10.2514/3.46062.
- [22] Marques, A. N., and Azevedo, J. L. F., “Numerical Calculation of Impulsive and Indicial Aerodynamic Responses Using Computational Aerodynamics Techniques,” *Journal of Aircraft*, Vol. 45, No. 4, 2008, pp. 1112–1135. doi:10.2514/1.32151.
- [23] Marques, A. N., and Azevedo, J. L. F., “A z Transform Discrete-Time State Space Formulation for Aeroelastic Stability Analysis,” *Journal of Aircraft*, Vol. 45, No. 5, 2008, pp. 1564–1578. doi:10.2514/1.32561.
- [24] Chen, P. C., Zhang, Z., and Eli, L., “Design-Oriented Computational Fluid Dynamics-Based Unsteady Aerodynamics for Flight-Vehicle Aeroelastic Shape Optimization,” *Journal of Aircraft*, Vol. 53, No. 12, 2015, pp. 3603–3619. doi:10.2514/1.J054024.

- [25] Opgenoord, M. M. J., Drela, M., and Willcox, K. E., “Physics-Based Low-Order Model for Transonic Flutter Prediction,” *AIAA Journal*, Vol. 56, No. 4, 2018, pp. 1519–1531. doi:10.2514/1.J056710.
- [26] Lucia, D. J., Beran, P. S., and Silva, W. A., “Reduced-order modeling: new approaches for computational physics,” *Progress in Aerospace Sciences*, Vol. 40, No. 1, 2004, pp. 51 – 117. doi:10.1016/j.paerosci.2003.12.001.
- [27] Benner, P., Gugercin, S., and Willcox, K., “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems,” *SIAM Review*, Vol. 57, No. 4, 2015, pp. 483–531. doi:10.1137/130932715.
- [28] Sirovich, L., “Turbulence and the Dynamics of Coherent Structures. Part 1: Coherent Structures,” *Quarterly of Applied Mathematics*, Vol. 45, No. 3, 1987, pp. 561–571. doi:10.1090/qam/910462.
- [29] Berkooz, G., Holmes, P., and Lumley, J. L., “The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows,” *Annual Review of Fluid Mechanics*, Vol. 25, 1993, pp. 539–575. doi:10.1146/annurev.fluid.25.1.539.
- [30] Dowell, E. H., and Hall, K. C., “Modeling of Fluid-Structure Interaction,” *Annual Review of Fluid Mechanics*, Vol. 33, No. 1, 2001, pp. 445–490. doi:10.1146/annurev.fluid.33.1.445.
- [31] Thomas, J. P., Dowell, E. H., and Hall, K. C., “Modeling Viscous Transonic Limit Cycle Oscillation Behavior using a Harmonic Balance Approach,” *Journal of Aircraft*, Vol. 41, No. 6, 2004, pp. 1266–1274. doi:10.2514/1.9839.
- [32] Silva, W. A., “Application of Nonlinear Systems Theory to Transonic Unsteady Aerodynamic Responses,” *Journal of Aircraft*, Vol. 30, No. 5, 1993, pp. 660–668. doi:10.2514/3.46395.
- [33] Missoum, S., Dribusch, C., and Beran, P., “Reliability-Based Design Optimization of Nonlinear Aeroelasticity Problems,” *Journal of Aircraft*, Vol. 47, No. 3, 2010, pp. 992–998. doi:10.2514/1.46665.
- [34] Lam, R., Allaire, D. L., and Willcox, K., “Multifidelity Optimization using Statistical Surrogate Modeling for Non-Hierarchical Information Sources,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA, 2015. doi:10.2514/6.2015-0143.
- [35] Poloczek, M., Wang, J., and Frazier, P., “Multi-Information Source Optimization,” *Advances in Neural Information Processing Systems 30*, Curran Associates, Inc., 2017, pp. 4291–4301.
- [36] Rasmussen, C. E., and Williams, C. K. I., *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.
- [37] Cover, T. M., and Thomas, J. A., *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*, Wiley-Interscience, 2006.
- [38] Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., and Richet, Y., “Fast Parallel Kriging-Based Stepwise Uncertainty Reduction With Application to the Identification of an Excursion Set,” *Technometrics*, Vol. 56, No. 4, 2014, pp. 455–465. doi:10.1080/00401706.2013.860918.

- [39] Isogai, K., “On the Transonic-dip Mechanism of Flutter of a Sweptback Wing,” *AIAA Journal*, Vol. 17, No. 7, 1979, pp. 793–795. doi:10.2514/3.61226.
- [40] Alonso, J. J., and Jameson, A., “Fully-Implicit Time-Marching Aeroelastic Solutions,” *32nd Aerospace Sciences Meeting and Exhibit*, AIAA, Reno, NV, 1994. doi:10.2514/6.1994-56.
- [41] Bisplinghoff, R. L., Ashley, H., and Halfman, R. L., *Aeroelasticity*, Dover Publications, Mineola, New York, 1996.
- [42] Dowell, E. H., *A Modern Course in Aeroelasticity*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2014. doi:10.1007/1-4020-2106-2.
- [43] Economon, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J., “SU2: An Open-Source Suite for Multiphysics Simulation and Design,” *AIAA Journal*, Vol. 54, No. 3, 2016, pp. 828–846. doi:10.2514/1.J053813.
- [44] Jacobson, K. E., Kiviaho, J. F., Kennedy, G. J., and Smith, M. J., “Evaluation of time-domain damping identification methods for flutter-constrained optimization,” *Journal of Fluids and Structures*, Vol. 87, 2019, pp. 174 – 188. doi:10.1016/j.jfluidstructs.2019.03.011.
- [45] Hall, K. C., Thomas, J. P., and Dowell, E. H., “Proper Orthogonal Decomposition Technique for Transonic Unsteady Aerodynamic Flows,” *AIAA Journal*, Vol. 38, No. 10, 2000, pp. 1853–1862. doi:10.2514/2.867.
- [46] Sanchez, R., Kline, H. L., Thomas, D., Variyar, A., Righi, M., Economon, T. D., Alonso, J. J., Palacios, R., Dimitriadis, G., and Terrapon, V., “Assessment of the fluid-structure interaction capabilities for aeronautical applications of the open-source solver SU2,” *VII European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2016)*, Crete Island, Greece, 2016. doi:10.7712/100016.1903.6597.
- [47] Hassig, H. J., “An Approximate True Damping Solution of the Flutter Equation by Determinant Iteration,” *Journal of Aircraft*, Vol. 8, No. 11, 1971, pp. 885–889. doi:10.2514/3.44311.
- [48] Johnson, E. H., and Venkayya, V. B., “Automated Structural Optimization System (ASTROS). Volume 1. Theoretical Manual,” Tech. rep., NORTHROP CORP HAWTHORNE CA AIRCRAFT DIV, 1988.
- [49] Timme, S., Marques, S., and Badcock, K. J., “Transonic Aeroelastic Stability Analysis Using a Kriging-Based Schur Complement Formulation,” *AIAA Journal*, Vol. 49, No. 6, 2011, pp. 1202–1213. doi:10.2514/1.j050975.
- [50] Burt, D., Rasmussen, C. E., and Van Der Wilk, M., “Rates of Convergence for Sparse Variational Gaussian Process Regression,” *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 97, edited by K. Chaudhuri and R. Salakhutdinov, PMLR, Long Beach, California, USA, 2019, pp. 862–871.