

A Formal Approach to Empirical Dynamic Model Optimization and Validation

Luis G. Crespo*

National Institute of Aerospace, Hampton, VA 23666 USA

Eugene A. Morelli[†], Sean P. Kenny[‡] and Daniel P. Giesy[§]

NASA Langley Research Center, Hampton, VA 23681 USA

A framework was developed for the optimization and validation of empirical dynamic models subject to an arbitrary set of validation criteria. The validation requirements imposed upon the model, which may involve several sets of input-output data and arbitrary specifications in time and frequency domains, are used to determine if model predictions are within admissible error limits. The parameters of the empirical model are estimated by finding the parameter realization for which the smallest of the margins of requirement compliance is as large as possible. The uncertainty in the value of this estimate is characterized by studying the set of model parameters yielding predictions that comply with all the requirements. Strategies are presented for bounding this set, studying its dependence on admissible prediction error set by the analyst, and evaluating the sensitivity of the model predictions to parameter variations. This information is instrumental in characterizing uncertainty models used for evaluating the dynamic model at operating conditions differing from those used for its identification and validation. A practical example based on the short period dynamics of the F-16 is used for illustration.

Keywords: uncertainty quantification, system identification, model validation, empirical dynamic modeling.

Nomenclature

- \mathcal{B} Bounding set of the validation domain
 \mathcal{C} Complement set operator

*Associate Research Fellow, AIAA Senior Member.

[†]Research Engineer, Dynamic Systems and Control Branch, MS 308, AIAA Associate Fellow.

[‡]Senior Research Engineer, Dynamic Systems and Control Branch, MS 308.

[§]Research Mathematician, Dynamic Systems and Control Branch, MS 308.

d	Number of samples
e	Error
ϵ	Admissible error
g	Requirement functions
J	Dimension reduction error
\bar{J}	Cost function used for the maximum likelihood
L	Likelihood function
m	Aspect vector
μ_r	Mean dimension reduction error
n	Number of parameters in θ
n_g	Number of requirement functions in g
θ	Parameters of the empirical model
$\hat{\theta}$	Maximal-margin parameter estimate
θ_c	Geometric center of the bounding set
$\bar{\theta}$	Maximum-likelihood parameter estimate
$\tilde{\theta}$	Critical parameter value (CPV)
ρ	Parametric safety margin (PSM)
R	Noise covariance matrix
\bar{R}	Maximum likelihood noise covariance matrix
t	Time
$\bar{\sigma}$	Standard deviation of θ
Σ	Parameter covariance
$u(t)$	Input to the empirical model
\mathcal{V}	Validation domain in θ -space
Vol	Volume operator
w	Worst-case requirement function
ω_r	Maximal dimension reduction error

I. Introduction

Consider the linear time invariant (LTI) model described by

$$\dot{x}(t) = A(\theta)x(t) + B(\theta)u(t) \quad (1)$$

$$y(t) = C(\theta)x(t) + D(\theta)u(t) \quad (2)$$

where $\theta \in \mathbb{R}^n$ is the vector of model parameters whose value is unknown, $t \in \mathbb{R}$ is time, $x(t) \in \mathbb{R}^{n_x}$ is the state, $x(0) = x_0$ is the initial condition, $u(t) \in \mathbb{R}^{n_u}$ is the input, $y(t) \in \mathbb{R}^{n_y}$ is the output, $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$, $C \in \mathbb{R}^{n_y \times n_x}$ and $D \in \mathbb{R}^{n_y \times n_u}$ are the system matrices. The output of the physical system modeled by Eqs. (1-2) is measured at discrete time instants t_i . These measurements are described by

$$z(t_i) = y(t_i) + \nu(t_i) \quad (3)$$

where $z(t_i) \in \mathbb{R}^{n_y}$ is the measured output for $i = 1, \dots, N$, and $\nu(t_i) \in \mathbb{R}^{n_y}$ is measurement noise. In aircraft dynamics, the parameters in θ usually are stability and control derivatives. System identification entails using the $u(t_i)$ and $z(t_i)$ corresponding to a carefully chosen

maneuver to estimate the value of θ , and thereby identify the dynamic model. For most aircraft system identification, it is assumed that the input u is measured without error. Thus, we will denote by $u(t_i)$ the measurement of $u(t)$ and by $z(t_i)$ the noisy measurement of $y(t)$. The input and output set $u(t_i)$ and $z(t_i)$ corresponding to the system identification experiment will be referred to as the system identification data set.

For traditional output-error parameter estimation, it is assumed that the measurement noise ν is a white zero-mean Gaussian process, and no other source of uncertainty is present. The unknown parameter vector θ is estimated by minimizing a weighted sum of squares of the difference between measured output z and model output y , in a maximum likelihood formulation. This formulation yields both the estimated parameter vector $\bar{\theta}$, and an ellipsoidal uncertainty set $\Delta_e \subseteq \mathbb{R}^n$ centered at $\bar{\theta}$, which is expected to contain the true value of θ with a specified level of confidence. This set is often approximated by Δ , a hyper-rectangle whose sides are the confidence intervals of the parameter estimates. Hence, the identified dynamic model consists of Eqs. (1-2), the estimated parameter $\bar{\theta}$, and the uncertainty set Δ .

The identified model is often validated using an independent flight test maneuver. The corresponding data set, to be called a validation data set, consists of the input $u_v(t_i) \in \mathbb{R}^{n_u}$ and the output $z_v(t_i) \in \mathbb{R}^{n_y}$, for $i = 1, \dots, N_v$. The validation of the identified model depends on its ability to predict z_v given u_v . In conventional aircraft system identification, model validation is done heuristically without mathematical formalism. The identified and validated empirical model, along with probabilistic descriptions of uncertainty based on Δ , are used in Monte Carlo analyses to study the probable range of performance for the physical system under different sets of operating conditions and/or different inputs.

This paper proposes a new approach to parameter estimation and dynamic model validation based on post-flight analysis. The framework proposed, which integrates multiple requirements corresponding to arbitrary performance metrics and/or several input-output data sets, is applicable to dynamic models with arbitrary functional forms, e.g., a nonlinear dynamic model with nonlinear parameterization. For example, one requirement may be based on the difference between the measured output used for system identification, z , and the predicted output $y(\theta, t, u)$ using the l_2 -norm of the difference as a performance metric, whereas another may be based on the maximum offset between the validation output z_v and the predicted output $y(\theta, t, u_v)$ using the l_∞ -norm of the difference as a performance metric. Any value of θ leading to sufficiently small differences is an admissible realization of the model parameters. The collection of all of these realizations constitute the validation domain. Hence, the system in Eqs. (1-2) evaluated at any member of the validation domain yields a prediction that meets both validation requirements. The size and shape of the validation domain characterizes the uncertainty in the parameters of the empirical model. In this setting, parameter estimation is performed by identifying the realization of the validation domain for which the smallest margin of requirement compliance is as large as possible. This represents an optimized dynamic model relative to the selected validation criteria. The validation domain plays a key role in formulating uncertainty models used by Monte Carlo analyses. Other uses include comparing flight results with ground-based predictions, calibrating computational fluid dynamics calculations, stability and control flight testing and flight envelope expansion, accident investigation, adaptive and learning control, and fault detection.

This paper presents techniques for calculating parameter estimates given several validation requirements, bounding the validation domain, studying its dependence on admissible

prediction error set by the analyst, and evaluating the sensitivity of the model predictions to parameter variations. This paper is organized as follows. Section II presents a summary of current practice in aircraft system identification and uncertainty analysis using the time-domain output-error method. Section III presents the proposed model validation technique. The global sensitivity analysis of the model predictions is covered in Section IV. Section V discusses practical implementation issues. Section VI illustrates in detail the scope of the methodology using a simple example. This is followed by Section VII, where a brief comparison between the legacy and proposed strategies is made. Finally, Section VIII closes the paper with a few concluding remarks.

II. Current Practice

The elements of the unknown parameter vector θ in Eqs. (1-2) can be estimated from measured data using an output-error formulation in the time domain. In this approach, θ is chosen to minimize the cost function

$$\bar{J}(\theta, u, z, R) = \frac{1}{2} \sum_{i=1}^N [z(t_i) - y(\theta, t_i, u)]^\top R^{-1} [z(t_i) - y(\theta, t_i, u)] + \frac{N}{2} \ln(|R|) \quad (4)$$

where $R \in \mathbb{R}^{n_y \times n_y}$ is a weighting matrix. If R is known, then only the term involving the summation in Eq. (4) is relevant for the optimization, with the elements of the θ vector being the only unknowns. The cost function is then the weighted sum of squared differences between the measured output z and the model output y for the given input u ,

$$e_2(\theta, u, z, R) = \frac{1}{2} \sum_{i=1}^N [z(t_i) - y(\theta, t_i, u)]^\top R^{-1} [z(t_i) - y(\theta, t_i, u)] \quad (5)$$

This formulation is called output error because the unknown parameters are chosen to minimize the weighted squared errors between the measured output and the model output.

The weighting matrix R in Eq. (4) is in general arbitrary, but if the noise ν is assumed to be a sample from a stationary, zero-mean, white Gaussian process, and R is chosen as the corresponding discrete noise covariance, $E[\nu(t_i)\nu(t_j)^\top] = R\delta_{ij}$, then the minimization of the cost in Eq. (4) is equivalent to the maximization of a Gaussian likelihood function, $L(\theta, u, z, R)$, defined as[1]

$$L(\theta, u, z, R) = \frac{1}{[(2\pi)^{n_y}|R|]^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2} \sum_{i=1}^N [z(t_i) - y(\theta, t_i, u)]^\top R^{-1} [z(t_i) - y(\theta, t_i, u)] \right\} \quad (6)$$

The cost function in Eq. (4) is the negative logarithm of the likelihood function given in Eq. (6), except for a constant term which has no bearing on the optimization. Maximizing the likelihood function is equivalent to minimizing its negative logarithm. Therefore, the *maximum likelihood estimate* $\bar{\theta}$ is found as

$$\bar{\theta} = \underset{\theta}{\operatorname{argmax}} L(\theta, u, z, R) = \underset{\theta}{\operatorname{argmin}} \bar{J}(\theta, u, z, R) \quad (7)$$

Maximum likelihood estimates have several desirable properties as the number of data points N gets large, such being asymptotically unbiased (estimate approaches the “true” value), and

having repeated estimates approach a Gaussian distribution with variances that approach the theoretical minimum.

Because of the nonlinearity of the cost function in Eq. (4), along with the generally coupled dynamics in Eqs. (1-2), the output-error cost function depends nonlinearly on the unknown parameter vector θ and also on the noise covariance matrix R . Consequently, a nonlinear optimization algorithm must be used to solve the problem. Typically, a relaxation method,¹ that combines the modified Newton-Raphson method (also called Gauss-Newton method) and the simplex method, can be used to minimize e_2 in Eq. (5), initially with R assumed equal to the identity matrix. This yields the estimate $\bar{\theta}$. The maximum likelihood noise covariance, \bar{R} , which minimizes the cost in Eq. (4) with respect to R for fixed $\bar{\theta}$, is given by the closed-form expression[1]

$$\bar{R} = \frac{1}{N} \sum_{i=1}^N [z(t_i) - y(\bar{\theta}, t_i, u)] [z(t_i) - y(\bar{\theta}, t_i, u)]^\top \quad (8)$$

The relaxation method alternates between the nonlinear optimization used to find $\bar{\theta}$ using Eq. (5) and the evaluation of \bar{R} using Eq. (8), until both estimates converge using selected convergence criteria.¹ Although there is no formal proof that the relaxation method will converge, extensive practical experience has shown that it does converge. This procedure has been mechanized in software and used successfully for a wide variety of aerospace dynamic modeling problems. Reference [1] provides further theoretical and practical details on the output-error parameter estimation method and the solution procedure outlined here. The output-error solution yields maximum likelihood estimates for both the unknown parameter vector θ and the noise covariance matrix R .

The procedures for calculating the parameter covariance matrix Σ , and an uncertainty for θ are presented next. The diagonal elements of Σ are the variances of the parameter estimates whereas the off-diagonal elements are the covariances. The Cramer-Rao lower bound for the covariance matrix, also called the dispersion matrix $D \in \mathbb{R}^{n \times n}$, is defined as

$$D = \left[\sum_{i=1}^N S^\top(t_i) \bar{R}^{-1} S(t_i) \right]^{-1} \quad (9)$$

where $S \in \mathbb{R}^{n_y \times n}$, called the output sensitivity matrix, quantifies the sensitivity of the model outputs to the model parameters at the maximum likelihood estimate:

$$S(t_i) = \left. \frac{\partial y(t_i)}{\partial \theta} \right|_{\theta=\bar{\theta}} \quad (10)$$

Output sensitivities can be calculated analytically by solving a set of sensitivity equations derived by differentiating Eqs. (1-2) with respect to θ , or by applying a finite difference technique to Eqs. (1-2), see Ref. [1] for details. The first-order approximation of the sensitivity of the model output to uncertainty is given by the product of S and the range of uncertainty. If the columns of the output sensitivity matrix S are non-zero and have low pairwise correlations, then each model parameter has a distinguishable effect on the output $y(t)$, and Eq. (7) is a well-conditioned optimization problem leading to an accurate parameter estimate with small uncertainties.

The connection between the Cramer-Rao bound D and the input u is complicated and nonlinear, even when the dynamic model is linear, as in Eqs (1, 2). Measured output noise levels (quantified by \bar{R}) also affect D . Roughly speaking, D is determined by squared noise to signal ratio, where \bar{R} represents the noise variance (square of the noise level), and the output sensitivities represent the signal strength for the parameter estimation problem. Because the covariance matrix Σ for maximum likelihood estimates asymptotically approaches the Cramer-Rao lower bound D as N increases, it is common practice to assume that the parameter covariance matrix Σ is approximately equal to D evaluated at the maximum likelihood estimate, i.e., $\Sigma \approx D$. Modern flight instrumentation systems collect flight data at 50 Hz or more, so that even relatively short maneuvers have a large number of data points N . This fact is used to justify the approximation $\Sigma \approx D$.

As the number of data points increases, maximum likelihood parameter estimates are asymptotically Gaussian with mean value equal to $\bar{\theta}$ and covariance equal to Σ . The corresponding probability density function is

$$f_{\theta}(\theta) = \frac{1}{(2\pi)^{\frac{n}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} (\theta - \bar{\theta})^{\top} \Sigma^{-1} (\theta - \bar{\theta})\right) \quad (11)$$

The hyper-ellipsoidal uncertainty set corresponding to the $(1 - \alpha)\%$ confidence level is then

$$\Delta_e(\bar{\theta}, \alpha, \Sigma) = \left\{ \theta : (\theta - \bar{\theta})^{\top} \Sigma^{-1} (\theta - \bar{\theta}) \leq \eta^2(\alpha) \right\} \quad (12)$$

where $\bar{\theta}$ is the center of the ellipsoid, and the value of $\eta(\alpha)$ can be determined from $1 - \alpha = \text{erf}(\eta/\sqrt{2})$. The values of θ for which the equality in Eq. (12) holds are level curves of the density function in Eq. (11). The normalized eigenvectors of Σ are the principal axes of the ellipsoid in Eq. (12), and the corresponding eigenvalues are the squared inverse of the corresponding semi-axis lengths. A common choice is to use the 95% confidence intervals, for which $\eta(0.05) \approx 2$. A confidence interval is an interval that includes the parameter of interest with a given level of confidence for repeated experiments. More specifically, if confidence intervals are constructed across many separate data analyses of different experiments, the proportion of such intervals that contain the true value of the parameter will match the confidence level. This is guaranteed asymptotically as more data is available.

The hyper-ellipsoidal uncertainty set in Eq. (12) can be roughly approximated by the hyper-rectangular uncertainty set

$$\Delta(\bar{\theta}, \alpha, \Sigma) = [\bar{\theta}_1 - \eta(\alpha)\bar{\sigma}_1, \bar{\theta}_1 + \eta(\alpha)\bar{\sigma}_1] \times \dots \times [\bar{\theta}_n - \eta(\alpha)\bar{\sigma}_n, \bar{\theta}_n + \eta(\alpha)\bar{\sigma}_n] \quad (13)$$

where $\bar{\sigma}_j$, $j = 1, \dots, n$, are the standard deviations of the estimated parameters. Each term on the right hand side of Eq. (13) corresponds to the $(1 - \alpha)\%$ confidence interval of a component of θ for the case in which the parameters are uncorrelated. Hyper-rectangles have the advantage that the range of an individual parameter is not affected by the ranges of the other parameters. This is in contrast to bounding sets with other geometries, such as ellipsoids. In general, none of the sets in Eqs. (12) and (13) contain each other.

The standard procedure in aircraft system identification for many years was to quote the maximum likelihood estimates of the parameters $\bar{\theta}$, with 95% confidence intervals computed based on D , and the Gaussian noise assumption. However, in practice, it was found that model residuals (difference between measured output and model output) are typically not

white Gaussian with constant variance, as assumed in the theory, but in fact are colored (time correlated) with nonuniform variance. This occurs because in practical problems, the model residuals include not just measurement noise, as assumed in the theory, but also model structure errors from effects not included in the model, such as nonlinearities, structural responses, and unsteady aerodynamic effects. The mismatch between theory and practice results in estimated parameter uncertainty bounds that are overly optimistic (too small) when using the standard procedure described above. Reference [2] provides some heuristic techniques to address this problem. References [1, 3, 4] provide a rigorous solution, validated using both simulation data and flight test data, which involves improving the \bar{R} estimate, and using it to compute a corrected D matrix. This solution has become the dominant approach in recent years, because standard errors computed in this way match the scatter in parameter estimates from repeated experiments at the same flight conditions. The necessary correction can be done as a post-processing of the residual sequences from an output-error maximum likelihood solution, because only the D matrix elements (and consequently the estimated parameter uncertainties) are affected, but not the parameter estimate $\hat{\theta}$.

Further details on the material in this section can be found in Ref. [1], which also includes MATLAB[®] software that implements the procedures described here, along with many other procedures and techniques related to aircraft system identification.

III. Parameter Estimation and Validation Domain Bounding

This section presents techniques for calculating a parameter estimate and bounding the validation domain. Recall that the dynamic model evaluated at any parameter realization of this domain yields predictions that satisfy the validation requirements. The validation domain is, in general, geometrically complex and multi-dimensional. These features often preclude its exact determination and analysis. Techniques for calculating inner bounding sets (subsets) and outer bounding sets (supersets) of this domain are presented below. While the inner bounding sets are guaranteed to consist only of validation points, the outer bounding sets are guaranteed to contain all the members of the validation domain. Some of the uses of these bounding sets are listed in Section VII. For purposes of this paper, these bounding sets are hyper-rectangles having edges that are parallel to the coordinate axes of θ -space.

The dependency of the dynamic model in Eqs. (1-2) on θ makes the stability and performance metrics derived from this model dependent on θ . The error measure in Eq. (5) and

$$e_{\infty}(\theta, u, z, R) = \max_{t_i} \left\{ [z(t_i) - y(t_i, \theta, u)]^{\top} R^{-1} [z(t_i) - y(t_i, \theta, u)] \right\} \quad (14)$$

which is an l_{∞} -norm, are examples of these metrics. Metrics like these will be used to prescribe a set of validation requirements. A dynamic model will be regarded as adequate when the corresponding model predictions comply with all the requirements. A requirement will be prescribed as an inequality constraint on the value of the error between the target behavior and the model prediction. If $e(\theta)$ denotes such an error, and ϵ is the largest admissible value of $e(\theta)$, this inequality can be described as

$$g(\theta, \epsilon) = e(\theta) - \epsilon \leq 0 \quad (15)$$

The dynamic model is compliant with the requirement if and only if the inequality is satisfied. An extension to the multiple requirement case can be easily made. The multiplicity of

requirements may stem from having multiple error forms or from having multiple data sets. For instance, the satisfaction of the vector inequality

$$g(\theta, \epsilon) = [e_2(\theta, u, z, R) - \epsilon_1, e_\infty(\theta, u_v, z_v, R) - \epsilon_2] \leq 0 \quad (16)$$

means that the square of the weighted l_2 difference between the predicted data and the measured data does not exceed ϵ_1 , and that the weighted l_∞ -norm of the difference between the predicted data and the validation data does not exceed ϵ_2 .

The vector inequality $g(\theta, \epsilon) \leq 0$, where $g \in \mathbb{R}^{n_g}$, partitions the space of model parameters θ into two regions - a region where all requirements are met, and a region where at least one of them is violated. The *validation domain* \mathcal{V} , or requirement-compliant domain, defined as

$$\mathcal{V}(\epsilon) \equiv \bigcap_{j=1}^{n_g} \{\theta : g_j(\theta, \epsilon_j) \leq 0\} \quad (17)$$

is comprised of all parameter realizations leading to predictions meeting all of the validation criteria. Each term on the right hand side of this equation defines the validation domain corresponding to an individual requirement. An alternative expression for the validation domain is $\mathcal{V}(\epsilon) = \{\theta : w(\theta, \epsilon) \leq 0\}$, where

$$w(\theta, \epsilon) \equiv \max_{1 \leq j \leq n_g} g_j(\theta, \epsilon_j) \quad (18)$$

is the *worst-case requirement function*. This formulation allows describing \mathcal{V} through a single inequality constraint. The complement of the validation domain, $C[\mathcal{V}]$, where $C[\cdot]$ denotes the set complement operator, will be called the *requirement-noncompliant* domain. Parameter realizations within this domain will violate at least one of the requirements.

The underlying logic supporting the calculation of the uncertainty set in Section II is to bound the “true” value of θ via confidence intervals, whereas the rationale used here is identifying the infinitely many realizations of θ for which the model predictions are adequate; e.g., the realizations of θ for which the offset between the measurements and the models’ prediction is within admissible error bounds.

This formulation can also be used to impose mathematical attributes, not necessarily described by the system identification or validation experiments, upon the dynamic model. For instance, a local stability requirement for the open-loop system can be imposed by including in g the requirement function $\max\{\Re[\lambda_1(\theta)], \dots, \Re[\lambda_{n_x}(\theta)]\} - \epsilon$, where $\lambda_k, k = 1, \dots, n_x$ are the eigenvalues of $A(\theta)$, $\Re[\cdot]$ is the real part operator, and $\epsilon = 0$. In the setting of Eq. (16), the parameter realizations within the validation domain yield locally stable dynamic models whose predictions adequately describe the measured outputs. Another requirement that often comes up in practice is that a particular model parameter should have a specific sign. If θ_a is to be negative, include $g_j(\theta, \epsilon_j) = \theta_a$ in the g vector. For a positive θ_a , use $g_j(\theta, \epsilon_j) = -\theta_a$.

The value of $-w(\theta, \epsilon)$ can be interpreted as a margin of requirement compliance associated with θ . In this setting, the parameter estimate $\hat{\theta}$, to be called the *maximal-margin estimate*, is defined as

$$\hat{\theta} \equiv \underset{\theta}{\operatorname{argmin}} \{w(\theta, \epsilon)\} \quad (19)$$

Therefore, the empirical model evaluated at $\hat{\theta}$ yields predictions that maximize the smallest margin of requirement compliance, i.e., $\hat{\theta}$ makes the value of the largest component of g is as

small as possible. If $w(\hat{\theta}, \epsilon) < 0$, then $\hat{\theta}$ is an element of the validation domain. Otherwise, $\mathcal{V}(\epsilon)$ is empty.

Techniques for calculating inner bounds require a non-empty requirement-noncompliant domain, whereas those for generating outer bounds require that \mathcal{V} be bounded. Figure 1 shows a sketch for $n = 2$, where the boundary of \mathcal{V} is shown as a solid black line. Two rectangular inner bounding sets, denoted as $\mathcal{B}_{\text{inner}}$ and $\tilde{\mathcal{B}}_{\text{inner}}$, and two outer bounding sets, denoted as $\mathcal{B}_{\text{outer}}$ and $\tilde{\mathcal{B}}_{\text{outer}}$, all centered at θ_c , are also shown. Note that $\tilde{\mathcal{B}}_{\text{inner}}$ and $\tilde{\mathcal{B}}_{\text{outer}}$ are, respectively, the minimum and maximum area rectangles bounding \mathcal{V} having θ_c as their center. $\mathcal{B}_{\text{inner}}$ and $\mathcal{B}_{\text{outer}}$ are the optimal squares.




Figure 1. Bounding sets of the validation domain

The m -scaled infinity norm is instrumental for calculating hyper-rectangular bounding sets. For a vector $m \in \mathbb{R}^n$ with positive components, the m -scaled infinity norm of $a \in \mathbb{R}^n$ is defined as $\|a\|_m^\infty = \max\{|a_k|/m_k\}$. The vector m , called the *aspect vector*, sets the orientation of the positive semi-diagonal of the hyper-rectangle. For example, the aspect vector of $\mathcal{B}_{\text{inner}}$ and $\mathcal{B}_{\text{outer}}$ in Figure 1 is $m = [1, 1]$, while $\tilde{\mathcal{B}}_{\text{inner}}$ and $\tilde{\mathcal{B}}_{\text{outer}}$ use an aspect vector that maximizes and minimizes the bounding set area, respectively. Reference [5] provides background information for the developments that follow.

A. Subsets of the Validation Domain

Formulations for calculating hyper-rectangular subsets of \mathcal{V} are presented below. While empirical models evaluated at any point from within such a subset will comply with the validation criteria, models evaluated outside this set may or may not satisfy these criteria. We present two formulations, one for a fixed aspect vector m and another which calculates an optimal aspect vector m . The optimal bounding set resulting from the latter formulation minimizes the volume of the uncovered portion of the validation domain.

1. Fixed Aspect Vector

Techniques for calculating the largest hyper-rectangle centered at θ_c with aspect vector m that is fully contained in \mathcal{V} are presented next. If $g(\theta, \epsilon) \leq 0$ is the set of requirements imposed upon the dynamic model, the hyper-rectangle

$$\mathcal{B}_{\text{inner}}(\theta_c, m, \rho_{\text{inner}}) = \{\theta : \theta_c - \rho_{\text{inner}}m \leq \theta \leq \theta_c + \rho_{\text{inner}}m\} \quad (20)$$

where $\theta_c \in \mathcal{V}$, with ρ_{inner} specified as

$$\rho_{\text{inner}}(\epsilon) = \min_{\theta} \{\|\theta - \theta_c\|_m^\infty : w(\theta, \epsilon) \geq 0\} \quad (21)$$

is a subset of \mathcal{V} . The calculation of a ρ , called the *parametric safety margin* (PSM), as in Eq. (21), entails solving a nonlinear optimization problem. A value of θ at which the optimum of Eq. (21) occurs, $\tilde{\theta}$, is called a *critical parameter value* (CPV). A CPV is a parameter realization where the bounding set and the requirement-noncompliant domain touch. The small yellow circles shown in Figure 1 show where the CPVs are. The components of g taking the value of zero at a CPV are the *critical requirements*. The PSM, each CPV, and the critical requirements are all functions of the geometry of \mathcal{V} , which in turn depends on the admissible error ϵ .

The values of the constants ϵ , θ_c , and m must be set before solving Eq. (21). The value of the components of admissible error vector ϵ should be set according to the nature and importance of each requirement. Some requirements allow for a somewhat subjective selection of ϵ (e.g., allowing for a 10% larger admissible error may still yield a sufficiently accurate prediction), whereas others may not (e.g., requiring the dynamic model to be locally stable demands an $\epsilon = 0$). For the subjective cases, the sensitivity of \mathcal{V} to ϵ can be explored by calculating the monotonically increasing and possibly discontinuous function $\rho_{\text{inner}}(\epsilon)$. Discontinuities in $\rho_{\text{inner}}(\epsilon)$ may arise when the $w(\theta, \epsilon)$ for a fixed ϵ has several local maxima. Additional information on $\rho_{\text{inner}}(\epsilon)$ is provided in Section III-C. The values of θ_c and m , which prescribe the geometry of the bounding sets, can be chosen in multiple ways. The value of θ_c should yield an accurate prediction, and the value of m_k should be made proportional to the level of uncertainty in the value of θ_k . A natural choice for the center of the bounding sets is $\theta_c = \hat{\theta}$. The conventional approach of Section II can be used to set $m = \bar{\sigma}$. A poor selection of θ_c and m may yield no solution for an inner bounding set (e.g., a case where $\theta_c = \bar{\theta}$ is outside the validation domain), or may yield unnecessarily loose bounding sets. The subjectivity in the selection of θ_c and m can be eliminated using the strategies proposed in the next section.

Imposing additional requirements or lowering the admissible error ϵ (or both) cannot increase ρ_{inner} . Rather, these actions will shrink the validation domain, and thus, the associated bounding sets. Note that the combination of several requirements, not necessarily stringent by themselves, may lead to an empty \mathcal{V} . So, it is possible to impose two requirements each leading to a satisfactorily large validation domain by itself, but whose intersection is comparatively small. This would imply that the sets of parameter values in which the model satisfies either of the two requirements is large, but the set in which both requirements are satisfied is small.

2. Optimal Aspect Vector

Techniques for calculating the hyper-rectangle centered at θ_c of largest volume within \mathcal{V} are presented next. In contrast to the fixed- m formulation, now the optimal value of m is determined in an outer optimization loop

$$\max_m \{\text{Vol}(\mathcal{B}_{\text{inner}})\} \quad (22)$$

The resulting bounding set is given by Eqs. (20) and (21), where

$$\tilde{m}_{\text{inner}}(\epsilon) = \operatorname{argmax}_m \left\{ \rho_{\text{inner}}^n(\epsilon) \prod_{k=1}^n m_k, \|m\| = 1, m > 0 \right\} \quad (23)$$

is the optimal aspect vector. The PSM and bounding box corresponding to \tilde{m}_{inner} will be denoted as $\tilde{\rho}_{\text{inner}}(\epsilon)$ and $\tilde{\mathcal{B}}_{\text{inner}}$, respectively. Note that finding \tilde{m}_{inner} requires solving a nested optimization problem where a search for the PSM for a fixed m occurs in the inner loop, while a search for the optimal aspect vector \tilde{m} occurs in the outer loop. As before, the resulting hyper-rectangle is an inner bounding set of \mathcal{V} . For the same θ_c , the bounding set resulting from this formulation is better than the fixed- m formulation in the sense that the volume of the validation domain not covered by the bounding set is the smallest. Figure 1 shows the optimal bounding set $\tilde{\mathcal{B}}_{\text{inner}}$ as a red rectangle, and the corresponding CPVs as yellow circles.

B. Supersets of the Validation Domain

In this section, techniques for finding supersets of \mathcal{V} are developed. The empirical model evaluated at any point outside this superset is requirement-noncompliant, whereas the model evaluated at points inside this superset may either be requirement-compliant or requirement-noncompliant. As before, two formulations are presented - one for a fixed m , and another for an optimal m . The optimal bounding set resulting from the latter formulation contains all the parameter realizations that comply with the requirements, with minimal excess, i.e., without extending unnecessarily over the requirement-noncompliant domain.

1. Fixed Aspect Vector

A formulation for calculating the smallest hyper-rectangle centered at θ_c with fixed aspect vector m that contains \mathcal{V} is presented next. If $g(\theta, \epsilon) \leq 0$ is the set of requirements imposed upon the dynamic model, the hyper-rectangle

$$\mathcal{B}_{\text{outer}}(\theta_c, m, \rho) = \{\theta : \theta_c - \rho_{\text{outer}}m \leq \theta \leq \theta_c + \rho_{\text{outer}}m\} \quad (24)$$

with ρ_{outer} is specified by

$$\rho_{\text{outer}}(\epsilon) = \max_{\theta} \{\|\theta - \theta_c\|_m^\infty : g(\theta, \epsilon) \leq 0\} \quad (25)$$

is a superset of \mathcal{V} . As before, the calculation of the PSM requires solving a nonlinear optimization problem. By construction, the family of empirical models resulting from varying θ over $C[\mathcal{B}_{\text{outer}}]$ will violate at least one of the requirements. The PSM, the CPV, and

critical requirements introduced in Section III-A can be naturally extended to the supersets considered here. Figure 1 shows the outer bounding set for $m = [1, 1]$, labeled $\mathcal{B}_{\text{outer}}$, as a blue rectangle, and the corresponding CPV as a yellow circle.

As before, the values of θ_c , m , and ϵ should be prescribed according to engineering judgement. Although the center of the inner bounding set θ_c must be a point within the validation domain (otherwise the bounding set would not be contained by \mathcal{V}), the center of the outer bounding set can be anywhere. Actually, if the validation domain is a disconnected set, the selection of a θ_c outside the validation domain can yield smaller bounding sets, e.g., the smallest rectangle that contains two non-intersecting circles of equal diameter has a center that is outside both circles.

As with the inner bounding sets, the sensitivity of \mathcal{V} to ϵ can be studied by calculating the monotonically increasing and possibly discontinuous function $\rho_{\text{outer}}(\epsilon)$. These discontinuities are the result of $w(\theta, \epsilon)$ having multiple local minima. Additional information on $\rho_{\text{outer}}(\epsilon)$ is provided in Section III-C.

2. Optimal Aspect Vector

Techniques for calculating the hyper-rectangle centered at θ_c of smallest volume that fully contains \mathcal{V} are presented next. In contrast to the fixed- m formulation, the optimal value of m is determined in an outer optimization loop

$$\min_m \{\text{Vol}(\mathcal{B}_{\text{outer}})\} \quad (26)$$

The resulting bounding set is given by Eqs. (24) and (25), where

$$\tilde{m}_{\text{outer}}(\epsilon) = \underset{m}{\text{argmin}} \left\{ \rho_{\text{outer}}^n(\epsilon) \prod_{k=1}^n m_k, \|m\| = 1, m > 0 \right\} \quad (27)$$

is the optimal aspect vector. The PSM and bounding set corresponding to \tilde{m}_{outer} will be denoted as $\tilde{\rho}_{\text{outer}}(\epsilon)$ and $\tilde{\mathcal{B}}_{\text{outer}}$, respectively. As before, finding \tilde{m}_{outer} requires solving a nested optimization problem where a search for the PSM for a fixed m occurs in the inner loop, and a search for the optimal aspect vector \tilde{m} occurs in the outer loop. For the same θ_c , the bounding set resulting from the free- m formulation is better than that of the fixed- m formulation in the sense that the volume of the requirement-noncompliant domain covered by the outer bounding set is the smallest. Figure 1 shows the optimal superset (green rectangle) and the corresponding CPVs (yellow circles). Note that multiple CPVs lay on the boundary of both inner and outer optimal bounding sets.

The free- m formulations for calculating bounding sets of \mathcal{V} can be trivially extended to the case where the center of the bounding set θ_c , along with m , are determined in the outer optimization loop. This practice leads to hyper-rectangular bounding sets where the offset between them and \mathcal{V} is minimal. This class of bounding sets will not be pursued here.

C. Dependency of the Bounding Sets on the Admissible Error

This section considers possible complexities of \mathcal{V} , the bounding sets, and the PSMs, as the admissible error ϵ prescribing the requirements changes. A simple example is used to

illustrate the case where the monotonically increasing function $\rho(\epsilon)$ corresponding to both inner and outer bounding sets of \mathcal{V} is discontinuous.

Figure 2 shows an error function $e(\theta)$, defined by the gray line, for a single unknown parameter θ . A piece-wise linear $e(\theta)$ has been chosen to facilitate the exposition. In this case the bounding sets are intervals centered at θ_c , which is the global minimum of the error function. Note that the inequality $g = w = e(\theta) - \epsilon \leq 0$, which solely defines the validation domain, is the singleton θ_c when $\epsilon = 0$. Larger values of ϵ yield validation domains with infinitely many points. The validation domains corresponding to a continuum of ϵ values is colored in green. The validation domain for ϵ_1 comprises two intervals, whereas the validation domain for ϵ_3 is a single interval. The colored lines in the left figure represent the endpoints of the inner bounding interval of \mathcal{V} (dashed red line) and the outer bounding interval of \mathcal{V} (dashed blue line). The half length of the intervals is the PSM.

The right figure shows the PSMs ρ_{inner} and ρ_{outer} as a function of ϵ . The local maxima and minima of $e(\theta)$ yield discontinuities in $\rho_{\text{inner}}(\epsilon)$ and $\rho_{\text{outer}}(\epsilon)$, respectively. As expected, $\rho_{\text{outer}} > \rho_{\text{inner}}$ for all ϵ . The functions $\rho_{\text{inner}}(\epsilon)$ and $\rho_{\text{outer}}(\epsilon)$ relate the maximum value of the worst-case error (abscissa) that occurs within a box centered at $\bar{\theta}$ with aspect vector m of a given size (ordinate). In this particular example, these functions relate the maximum value of $e(\theta)$ that occurs within an interval centered at $\bar{\theta}$ having a half length equal to the PSM. The discontinuities in the PSM are cases where a differential change in the admissible error ϵ yields a significant change in the validation domain, and thus, in the bounding sets. In the case shown, each discontinuity is associated with $\mathcal{V}(\epsilon)$ changing between being simply-connected and being not simply-connected.




Figure 2. PSM of inner and outer bounding sets of \mathcal{V} as a function of ϵ

The following section examines the global sensitivity of all the components of g , which may well include metrics based on the output $y(t)$, to the model parameters as their values are free to range within the validation domain. The understanding of this section is not necessary for the remainder of this paper, and as such, it may be considered optional.

IV. Global Sensitivities

The output sensitivity matrix S in Eq. (10) and the parameter covariance matrix Σ in Eq. (9) provide valuable information on the sensitivity of the predicted output $y(t)$ to variation in the model parameters. If a particular parameter sensitivity is poor, then that parameter is dropped or fixed as part of the standard procedure in the system identification problem. This well-established approach is based on the local sensitivity of the identified model output $y(t)$ to the system identification input at the single parameter point $\bar{\theta}$. Performance metrics and requirements that are not based on $y(t)$ will likely exhibit different sensitivities, e.g., the sensitivity of the first natural frequency of the identified model to θ . Moreover, derivatives about other θ points cannot be expected to be in agreement with those at $\bar{\theta}$. In general, there is no basis to expect that the results of a global sensitivity analysis will be in agreement with those of a local sensitivity analysis, unless the range of variation in θ is small, or the figures of merit prescribing all the components of g are a weakly nonlinear function of θ .

Up to this point, the problem of assigning values to θ such that the system model satisfies all requirements has been considered. The acceptability of candidate values of θ is based on an error metric which measures, for instance, how closely the model output matches the test results when the candidate value is used for the model parameters. In this section, methods are explored for determining if, from the point of view of matching test results, the model is so insensitive to the value assigned to specific components of θ that, for analysis purposes, those components may be assigned fixed values without materially changing the quality of the prediction. If such insensitive parameters can be identified, the dimensionality of the parameter identification problem is effectively reduced. Similar ideas are used in Ref. [6]. A mathematical formalism to describe these notions is presented next.

The effect that the value of a subset of the model's parameters has on the quality of the model's predictions is studied next. To this end, we partition the parameter vector θ into two subvectors, the vector f of parameters whose *range of variation* within \mathcal{V} will be preserved, and the vector r of parameters whose *value* within \mathcal{V} will be kept fixed. This partition will be represented as $\theta = [f, r]$. Denote by \mathcal{F} and \mathcal{R} the projections of the validation domain onto the f - and r -subspace respectively

$$\mathcal{F} = \{f : \exists r \text{ for which } [f, r] \in \mathcal{V}\} \quad (28)$$

$$\mathcal{R} = \{r : \exists f \text{ for which } [f, r] \in \mathcal{V}\} \quad (29)$$

The predictions made by the dynamic model are insensitive to r when $g([f, r], \epsilon) \approx g([f, \hat{r}], \epsilon)$ for all $[f, r] \in \mathcal{V}$, where \hat{r} is any fixed value of r in \mathcal{R} . A metric for evaluating the approximation error incurred by using $g([f, \hat{r}], \epsilon)$ instead of $g([f, r], \epsilon)$ is the *dimension reduction error* $J(\hat{r})$, defined as

$$J(\hat{r}) \equiv \int_{\mathcal{V}} \int [g([f, r], \epsilon) - g([f, \hat{r}], \epsilon)]^T Q [g([f, r], \epsilon) - g([f, \hat{r}], \epsilon)] df dr \quad (30)$$

where Q is a positive definite matrix. The matrix Q can be used to scale the components of g according to the units of the error e prescribing it, and to emphasize particular requirements. A small value of $J(\hat{r})$ for all $\hat{r} \in \mathcal{R}$ indicates global insensitivity of g , thus of e and y , to the value that r might take in \mathcal{R} .

A couple of metrics for evaluating key features of the function $J(\hat{r})$ as \hat{r} ranges over \mathcal{R} are introduced next. The first of these metrics the *maximal dimension reduction error*, ω_r ,

given by

$$\omega_r = \max_{\hat{r}} \left\{ \frac{J(\hat{r})}{\text{Vol}(\mathcal{R})} : \hat{r} \in \mathcal{R} \right\} \quad (31)$$

The metric ω_r is a tight upper bound on all possible dimension reduction errors that may occur by fixing the value of \hat{r} in \mathcal{R} . The second metric is the *mean dimension reduction error*, μ_r , given by

$$\mu_r = \int_{\mathcal{R}} \frac{J(\hat{r})}{\text{Vol}(\mathcal{R})} d\hat{r} \quad (32)$$

While ω_r depends on the extreme and possibly unlikely case at which $J(\hat{r})$ attains its maximum, μ_r is the average value of $J(\hat{r})$ when \hat{r} is uniformly distributed over \mathcal{R} . The sensitivity metrics ω_r and μ_r can be used to rank the model's parameters and determining which of them can be fixed at a constant value. The best of such values is the point where J attains its minimum. Note that rankings based on ω_r and μ_r may not agree. Further notice that fixing the value of the parameters in \hat{r} effectively reduces the dimension of the unknown parameter space from n to $\dim(f)$.

The sensitivity metrics ω_r and μ_r can be evaluated using bounding sets of \mathcal{V} . Details on how to evaluate them numerically are provided next. Let $\mathcal{B}(\bar{\theta}, \hat{m}, \hat{\rho})$ be a superset of $\mathcal{V}(\epsilon)$. The partition $\theta = [f, r]$ leads to $\mathcal{B} = \mathcal{B}^f \times \mathcal{B}^r$, where $\mathcal{B}^f = \{f : \bar{f} - \hat{\rho}\hat{m}^f \leq f \leq \bar{f} + \hat{\rho}\hat{m}^f\}$, $\mathcal{B}^r = \{r : \bar{r} - \hat{\rho}\hat{m}^r \leq r \leq \bar{r} + \hat{\rho}\hat{m}^r\}$, $\hat{m} = [\hat{m}^f, \hat{m}^r]$, and $\theta_c = [f_c, r_c]$. In this setting, the sensitivities in Eq. (31) and Eq. (32) are given by

$$\omega_r = \frac{1}{\text{Vol}(\mathcal{R})} \max_{\hat{r} \in \mathcal{B}^r} \left\{ \int_{\mathcal{B}^f} \int_{\mathcal{B}^r} h(f, r, \hat{r}) I_{\mathcal{V}}([f, r]) df dr \right\} \quad (33)$$

$$\mu_r = \frac{1}{\text{Vol}(\mathcal{R})} \int_{\mathcal{B}^r} \int_{\mathcal{B}^f} h(f, r, \hat{r}) I_{\mathcal{V}}([f, r]) df dr \quad (34)$$

where $h(f, r, \hat{r})$ is the integrand of Eq. (30), I is the indicator function, and

$$\text{Vol}(\mathcal{R}) = \int_{\mathcal{B}^r} I_{\mathcal{R}}(r) dr \quad (35)$$

The indicator function $I_{\mathcal{A}}(a)$ is equal to one when $a \in \mathcal{A}$, otherwise it is equal to zero. The integrals above can be evaluated using sampling. This practice leads to approximation errors which diminish only as $O(1/\sqrt{d})$, where d is the number of samples. This results in the usual diminished returns of ordinary statistical estimators: to halve the error in the estimate of ω_r we must quadruple d . While the integrals can also be evaluated using Riemann's sums, the computational cost, which grows exponentially with n , restricts this approach to problems with small number of uncertainties. As expected, the sets resulting from the optimal- m formulations, which bound more tightly the validation domain, yield, for a given number of samples, more accurate estimates of the integrals, thus, of the sensitivity metrics.

The sensitivity of the model prediction to a particular model's parameter prescribes the effort we should devote to modeling the uncertainty on its value. Besides providing insight into the characterization of uncertainty, sensitivity information also enables establishing a basic notion of causality. By knowing which parameters are important, better system identification experiments can be designed and critical requirements can be tightened or relaxed

accordingly. Furthermore, if the sensitivity metric for a particular parameter is close to zero, the corresponding component of the system matrices in Eqs. (1-2) can be fixed in advance and the validation domain of the new model, now in the $\dim(f)$ -dimensional subspace, can be studied. If the dynamic model is under-parameterized the validation domain corresponding to acceptable values of the admissible error ϵ will be empty. When the validation domain is non-empty, the model can be either overly-parameterized or adequately parameterized. The model is overly-parameterized when the value of some parameters can be fixed and the model predictions corresponding to the validation domain associated with the remaining parameters are about as good as the original. The model is adequately-parameterized when fixing the value of any of the parameters yields a substantially degraded model prediction.

V. Practical Implementation

Guidelines for practical use of the framework proposed are given next. The following step-by-step procedure provides a high-level description of the tasks required to estimate the model's parameters, characterize the uncertainty in their value, and validate the empirical model.

1. Chose the structure of the dynamic model according to the physics governing the system, or identify an adequate model structure from the data, using model structure determination techniques.¹ The dependence of the model output on the states, controls, and parameter vector θ can be arbitrary.
2. Qualitatively prescribe the set of requirements to be imposed upon the model. Stability and performance requirements in the time and frequency domains, as well as constraints on the allowable range of variation for the parameter vector θ , can be included. These requirements can bound the offset between the observations and model predictions, or can force the model to satisfy desired mathematical properties.
3. Prescribe the above requirements quantitatively. This task entails creating a g for each of the requirements such that $g \leq 0$ implies requirement satisfaction whereas $g > 0$ implies its violation. Requirements bounding the offset between the observations and the model predictions can be cast as $g = e - \epsilon$, where the error e is any norm of the offset and ϵ is the admissible error. Time and frequency domain specifications can be prescribed by setting allowable ranges of variation for key figures of merit, e.g., if $t_s(\theta)$ is the settling time to a step input and 1 is the maximum allowable settling time, use $g = t_s(\theta) - 1$.
4. Scale the components of g to make their range of variation comparable. Keep in mind the requirement functions have different units and describe different performance metrics. Whereas scaling affects the parameter estimate $\hat{\theta}$, it does not affect the validation domain bounds since their calculation only depends on the sign of g .
5. Find the maximal-margin parameter estimate $\hat{\theta}$ in Eq. (19). The best fixed-parameter empirical model is given by the structure set in Step (1) evaluated at $\hat{\theta}$.
6. Find a parameter realization for which the empirical model satisfies all the requirements. Denote this realization as θ^* . θ^* can be searched for by using sampling, or can

be made equal to the maximal-margin parameter estimate provided that $w(\hat{\theta}, \epsilon) < 0$. Failing to find a value of θ for which this inequality holds may indicate that the model structure does not capture the problem physics well, or that the requirements are overly stringent. If this is the case, revisit steps 1 through 3 to mend these deficiencies. The existence of θ^* indicates that the validation domain is not empty.

7. Calculate inner bounding sets of the validation domain using either Eq. (21) or Eq. (23) for $\theta_c = \theta^*$. The latter formulation yields tighter sets.
8. Calculate outer bounding sets of the validation domain using either Eq. (25) or Eq. (27) for $\theta_c = \theta^*$. The latter formulation yields tighter sets.
9. The family of validated dynamic models is given by the model structure prescribed in Step 1, and the bounds on the validation domain in Steps (7) and (8). While all the members inside the inner bounding set yield validated models, all the members outside the outer bounding set yield invalidated models.
10. For a global sensitivity analysis use the developments of Section IV. This analysis enables (i) ranking the parameters in θ according to the extent which they affect the model's ability to meet the requirements, and (ii) determining which parameters can be assumed to take on a fixed value without significantly degrading the accuracy of the model prediction.
11. Use the outer bounding set of the validation domain as the support set of uncertainty models for θ . This model is commonly used to study the variability on the model predictions caused by parameter uncertainty using Monte Carlo analysis. The larger the offset between the validation domain and its outer bounding set, the larger the conservatism in the analysis.

A few remarks on the computational requirements of the above procedure are in order. Note that the optimization problems in Eqs. (7, 19, 21, 23, 25, 27), are nonlinear in general. These problems can be solved by most of the nonlinear programming algorithms available in numerical computation software. For the unconstrained problems in Eqs. (7, 19, 23, 27), techniques such as Nelder-Mead (nonlinear simplex), trust-region, and quasi-Newton algorithms are applicable. For the constrained optimization problems in Eqs. (21, 25), sequential quadratic programming, interior-point, active-set, and trust-region reflective, can be used. In this work, `fminsearch`, `fminunc` and `fmincon` from MATLAB[®] were used.

In any non-convex optimization problem there is always the possibility of convergence to a local optimum instead of the global optimum. While convergence to a non-global optimum in the unconstrained problems of Eqs. (7, 19, 23, 27), is not critical, this is not the case for the constrained cases in Eqs. (21, 25). When this anomaly occurs in the non-critical cases, the resulting sets are suboptimal bounding sets, e.g., there exists a larger inner bounding set. When this anomaly occurs in the critical cases, the resulting set is not an actual bounding set. Absolute guarantees of convergence to the global optimum are not possible, but a variety of algorithmic safeguards can be used to compensate for this deficiency. For instance, in the calculation of an outer bounding set, g can be evaluated at few sample points outside the resulting set. If $g \leq 0$ for any of these sample points, which will indicate that the calculated set is flawed, any of those sample points can be used as initial condition in subsequent

searches. Note that Monte Carlo analysis cannot be used to calculate bounding sets, or to prove that any set in the parameter space is indeed a bounding set. This is the result of not being able to assess the regions of the parameter space where there are no sample points. For instance, if the values w at selected sample points within the set $\mathcal{B}_{\text{inner}}$ are all negative, the possibility of w being positive in between any pair of sample points still exists. While finding one sample point in \mathcal{B} for which $w > 0$ is enough to prove that \mathcal{B} is not an inner bounding set of \mathcal{V} , finding that $w < 0$ at any finite number of sample points is not enough to establish that \mathcal{B} is contained in \mathcal{V} .

The computational effort of calculating bounding sets depends on the effort required to evaluate g , the number of uncertain parameters n , and the convergence properties of the nonlinear optimizer. This calculation may entail performing time simulations, calculating eigenvalues and/or frequency domain metrics, among many other possibilities. A detailed numerical analysis of the convergence to the optimum as a function of these parameters is beyond the scope of this paper. However numerical experiments based on the example below take a few hundred function evaluations of g , and about 4 minutes (2×3.2 GHz Quad-Core) to converge. Because this technology is to be applied post-flight, and not in real time, the CPU time associated with these calculations does not play a critical role. A concrete example, where the main ideas above are developed in detailed, is presented next.

VI. Example

An example case using a dynamic model with 6 unknown parameters is studied next. An LTI model for the short-period dynamics was extracted from the nonlinear F-16 simulation in [1] using central finite differences. A steady level flight condition for $\alpha = 5$ deg was chosen. The model structure is the standard form given in Eqs. (1-2) where u is the deflection of the stabilator, $x = [\alpha, q]^\top$, α is the angle of attack, q is the pitch rate, and $y = x$. The system matrices are

$$A(\theta) = \begin{bmatrix} \theta_1 & \theta_2 \\ \theta_4 & \theta_5 \end{bmatrix}, B(\theta) = \begin{bmatrix} \theta_3 \\ \theta_6 \end{bmatrix}, C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, D = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (36)$$

The system identification data and the validation data were generated using the nonlinear simulation with Gaussian noise for two different doublet input sequences. As a result, the system identification input $u(t_i)$ and output $z(t_i)$, as well as the validation input $u_v(t_i)$ and output $z_v(t_i)$, are available. Following the standard approach of Section II, we calculate the maximum likelihood estimate $\bar{\theta} = [-0.6454, 0.9066, -0.1538, -3.7948, -1.2015, -6.5242]$, and the standard deviation $\bar{\sigma} = [0.0214, 0.0104, 0.0208, 0.0329, 0.0265, 0.0682]$. This information fully prescribes the uncertainty set $\Delta(\bar{\theta}, 0.05)$ in Eq. (13).

Validation requirements based on the error in Eq. (5) will be imposed upon the dynamic model. The first requirement is $g_1 = e_s - \epsilon$, where

$$e_s(\theta) = \frac{e_2(\theta, u, z)}{e_2(\bar{\theta}, u, z)} \quad (37)$$

$\epsilon \geq 1$, and u and z correspond to the system identification data set for $R = \bar{R}$. This requirement implies that the error incurred by θ must not exceed the minimum value by more than $100(\epsilon - 1)$ %. Note that the function $e_s(\theta)$, which depends upon six unknown parameters, has a minimum equal to 1 at $\theta = \bar{\theta}$. The normalization in (37) enables the

comparison of the e_2 error to the value attained by the maximum likelihood estimate. The left subplots in Figure 3 show 2-dimensional contours of $e_s(\theta)$ within the 95% confidence box $\Delta(\bar{\theta}, 0.05)$. Note that the geometric center of the box is both the maximum likelihood




Figure 3. Contours of the error function for one (left) and two requirements (right).

estimate $\bar{\theta}$ (marked with a + sign) and the maximal-margin parameter estimate $\hat{\theta}$ (marked with a \times sign). Elements of θ that are not varied in the figure are fixed at their maximum likelihood value. The largest error caused by uncertainty in θ_5 and θ_6 within the rectangle pictured in the upper left is 27% larger than the error at $\bar{\theta}$, while the largest error caused by uncertainty in θ_1 and θ_2 is 11% larger. Further notice that the contours depict the boundary of the validation domain corresponding to different values of ϵ . Hence, the validation domains have an ellipsoid-like shape centered at $\bar{\theta}$ whose principal axes are fairly independent of ϵ . The boundary of $\mathcal{V}(1.025)$ is shown as a black line in all subplots.

Now consider two requirements, the one above and one based on the validation data set. The constraint corresponding to a second requirement is $g_2 = e_v - \epsilon$, where

$$e_v(\theta) = \frac{e_2(\theta, u_v, z_v)}{e_2(\theta_v, u_v, z_v)} \quad (38)$$

$\theta_v \equiv \min_{\theta} \{e_2(\theta, u_v, z_v)\}$, and u_v and z_v correspond to the validation data set. Note that θ_v is the maximum likelihood estimate corresponding to this data set. In general, the value of θ_v does not coincide with $\bar{\theta}$. The intent and structure of both requirements is the same. Furthermore, the same admissible error ϵ will be used in both constraints. As with e_s , the

minimum of e_v is also 1. This requirement implies that the variation in the error based on the validation data set incurred by θ must not exceed the minimum attainable value by more than $100(\epsilon - 1) \%$.

Figure 4 shows the measured and predicted angle of attack corresponding to the system identification data set (top) and the validation data set (bottom). The figure shows the measured output (blue line), which is contaminated by noise, and the model predictions for 3 values of θ . The predictions corresponding to $\bar{\theta}$ and θ_v , for which the prediction for each requirement is optimal, are shown in red and cyan, while those corresponding to θ_a , a value within $\Delta(\bar{\theta}, 0.05)$, is shown in black. The difference between the model predictions is barely visible. Figures 5 and 6 are used to better show such difference. Figure 5 shows




Figure 4. Input, measured output and model predictions for 3 sets of model parameters.

the unfiltered and filtered versions of the error between the measured output and the model prediction for $\bar{\theta}$ (red-dashed line) and θ_a (black-solid line). The filtered versions result from using a low-pass filter to remove the high frequency noise from the measurements. The error functions corresponding to θ_v , which are practically indistinguishable from those of $\bar{\theta}$, are omitted. Figure 6 shows the cumulative values of e_s and e_v for $\bar{\theta}$, θ_v and θ_a . By cumulative value we mean that the final integration time of the numerator in Eq. (37) and Eq. (38) is varied from 0 sec to 10 sec. Note that the values of e_s and e_v are equal to the terminal value of the cumulative errors. Further notice that dynamic models that describe well one of the measured outputs may be inaccurate predictors of the other one. For example, the dynamic model evaluated at θ_a exceeds the minimum matching errors of the system identification and validation data sets by 70%, and 285%, respectively. If we set 100% as the maximum allowable percentage of exceedance for both requirements, the dynamic model associated with θ_a will be outside of the validation domain even though its prediction of the system identification output is sufficiently accurate. Note also that even though $\bar{\theta}$ and θ_v

yield practically identical predictions, $e_s(\bar{\theta}) < e_s(\theta_v)$ and $e_v(\bar{\theta}) > e_v(\theta_v)$. The inequality




Figure 5. Error difference between measurement and prediction for 2 sets of model parameters.

constraint describing the satisfaction of both requirements for the same admissible error is $g = [g_1, g_2] = [e_s(\theta) - \epsilon, e_v(\theta) - \epsilon] \leq 0$. The parameter realizations satisfying both of these inequalities comprise the validation domain for both requirements. The boundary of $\mathcal{V}(\epsilon)$ can be displayed by plotting contours of the function $e_{sv}(\theta) = \max\{e_s(\theta), e_v(\theta)\}$, which is the worst-case requirement function in Eq. (18).

The right subplots in Figure 3 show 2-dimensional contours of $e_{sv}(\theta)$ within the 95% confidence box $\Delta(\bar{\theta}, 0.05)$. In this setting the maximal-margin parameter estimate, which is the parameter realization minimizing e_{sv} marked with a \times sign, differs from the maximum likelihood estimate $\bar{\theta}$, which is the parameter realization minimizing e_s marked with a $+$ sign. Notice that there are values of ϵ for which $\bar{\theta}$ is outside the validation domain. This is the case when $\epsilon = 1.025$, which is the value leading to the validation domain boundary shown in black. This implies that the parameter estimate leading to the best value of e_s (i.e., best estimate regarding the first requirement) yields a value of e_v that exceeds admissible levels. The maximal-margin parameter estimate, which is not even at the geometric center of $\mathcal{V}(1.025)$, is $\theta_c = [-0.6280, 0.9060, -0.1587, -3.7862, -1.2212, -6.5461]$. As expected, the minimum of $e_{sv}(\theta_5, \theta_6)$ and of $e_{sv}(\theta_1, \theta_2)$ are larger than the minimum of both $e_s(\theta_5, \theta_6)$ and $e_v(\theta_1, \theta_2)$. Values of ϵ below such a minimum yield an empty validation domain, i.e., there are no parameter realizations satisfying both requirements. The comparison between the left and right subplots of Figure 3 shows that the imposition of an additional requirement substantially increases the error spread within Δ . In contrast to the single requirement case shown to the left, the validation domains corresponding to small values of ϵ are no longer ellipsoids. Furthermore, note that for large values of ϵ , the principal axes of the family of ellipsoid-like curves shown to the right slightly differ from those shown to the left.




Figure 6. Cumulative errors associated with model predictions for 3 sets of model parameters.

These observations illustrate that the imposition of additional requirements, even in this simple case for which the structure of both requirements uses the same error used by the maximum-likelihood approach, can substantially affect both the maximal-margin parameter estimate and the geometry of the validation domain. Figure 3 illustrates how the imposition of an additional requirement reduces the size of the validation domain while ensuring that \mathcal{V} corresponding to a given set of m requirements is contained by the \mathcal{V} corresponding to any subset of $m - 1$ requirements out of these m .

Next, we calculate the PSM function $\rho_{\text{inner}}(\epsilon)$ for the center $\theta_c = \hat{\theta}$, and the fixed aspect vector $m = \bar{\sigma}/\|\bar{\sigma}\|$. Recall that the validation domain being bounded is a set in the 6-dimensional θ space. As before, we will consider a case with a single requirement based on Eq.(37), and another case with both requirements based on Eqs. (37) and (38). Figure 7 shows the ratio $\rho_{\text{inner}}(\epsilon)/\|\sigma\|$ for the one requirement case (red dashed line) and the two requirement case (blue solid line). For a given inner bounding box one must accept greater error in order to satisfy two requirements vs. only one requirement. For example, the largest error within the box $\Delta(\hat{\theta}, \alpha)$, which corresponds to the ordinate value 2 in Figure 7, is about 75% larger than the smallest attainable error. In contrast, the two requirement case leads to maximum errors that are about 265% larger than the minimum error. As it was observed in the 2-dimensional plots above, the imposition of the second requirement considerably increases the prediction error within the box. Note that the value of ϵ where $\rho_{\text{inner}}(\epsilon) = 0$ is 1.07 in the two requirement case and 1.0 in the one requirement case. This is consistent with the increase of the minimum of e_{sv} mentioned above. In this case the second requirement is the worst-case requirement for all values of ϵ (the PSM function corresponding to the second




Figure 7. $\rho_{\text{inner}}(\epsilon)/\|\bar{\sigma}\|$ for the one and two requirement cases having $m = \bar{\sigma}$.

requirement alone will lead to the same blue line). In general, the individual requirement that prescribes the worst-case requirement depends on the value of ϵ .

Next, we calculate inner bounds of $\mathcal{V}(\epsilon)$ for the two requirement case using the optimal- m formulation. Figure 8 shows the corresponding PSM, ρ_{inner} . Recall that the aspect vector \tilde{m}_{inner} maximizes the volume of the inner bounding set. In this case the optimal \tilde{m}_{inner} changes slightly with ϵ . While the aspect vector used in Figure 7 is $m = \bar{\sigma}/\|\bar{\sigma}\| = [0.24, 0.12, 0.24, 0.38, 0.30, 0.79]$, the optimal ones are closely scattered about $m = [0.36, 0.12, 0.30, 0.55, 0.18, 0.64]$.

Next, we calculate outer bounds of $\mathcal{V}(\epsilon)$ for the two requirement case using the optimal- m formulation. Figure 8 shows the corresponding PSM, ρ_{outer} . Recall that the aspect vector \tilde{m}_{outer} minimizes the volume of the outer bounding set. In contrast to the previous case, the optimal \tilde{m}_{outer} changes considerably with ϵ . Figure 8 indicates that the volume of the outer bounding set is several orders of magnitude larger than the volume of the inner bounding set. This difference indicates that the maximum likelihood estimate $\bar{\theta}$ is either far from the centroid of $\mathcal{V}(\epsilon)$ (as the left plots of Figure 3 suggest) or that $C[\mathcal{V}]$ has long spikes intruding into the validation domain.

Finally, the sensitivity analysis of Section IV is used to rank the six parameters in θ . To this end we will use the bounding set of $\mathcal{V}(1.15)$. Therefore, prediction errors that exceed in up to 15% the minimum attainable e_2 error for both the system identification data set and validation data set are admissible. The optimal PSM, optimal aspect vector, and volume of the inner bounding set are $\tilde{\rho}_{\text{inner}} = 0.0332$, $\tilde{m}_{\text{inner}} = [0.3914, 0.1048, 0.2902, 0.4574, 0.2120, 0.7052]$ and $\text{Vol}(\mathcal{B}_{\text{inner}}) = 7.0154 \times 10^{-11}$, while the values corresponding to the outer bounding set are $\tilde{\rho}_{\text{outer}} = 0.3261$, $\tilde{m}_{\text{outer}} = [0.2204, 0.3945, 0.2554, 0.7975, 0.2198, 0.2147]$, and $\text{Vol}(\mathcal{B}_{\text{outer}}) = 0.06445$. The dimension reduction error, $J(\hat{r})/\text{Vol}(\mathcal{R})$ for $\hat{r} = \theta_1, \dots, \theta_6$, corresponding to $\mathcal{B}_{\text{outer}}$ are shown in Figure 9. Vertical lines at the values of the maximum likelihood estimate are superimposed. The value of $J(\hat{r})/\text{Vol}(\mathcal{R})$ at the maximum likelihood estimate as well as the derivative at such a point, cannot be used to describe global




Figure 8. PSMs and volumes of the optimal- m bounding sets as a function of the admissible error ϵ .

sensitivities. The sensitivities ω_r and μ_r can be readily evaluated from these functions. Sensitivities based on the maximal dimension reduction error are $\omega_{\theta_1} = 0.0161$, $\omega_{\theta_2} = 0.1729$, $\omega_{\theta_3} = 0.0157$, $\omega_{\theta_4} = 0.0229$, $\omega_{\theta_5} = 0.0432$, and $\omega_{\theta_6} = 0.0013$. This yields the parameter ranking in decreasing order of importance [2, 5, 4, 1, 3, 6]. Conversely, sensitivities based on the mean dimension reduction error are $\mu_{\theta_1} = 0.0033$, $\mu_{\theta_2} = 0.0475$, $\mu_{\theta_3} = 0.0038$, $\mu_{\theta_4} = 0.0054$, $\mu_{\theta_5} = 0.0086$, and $\mu_{\theta_6} = 0.0003$. The parameter ranking according to μ is [2, 5, 4, 3, 1, 6]. In both cases there is a two orders of magnitude difference between the most and the least important parameters. The functions shown in Figure 9 illustrate how J , thus the sensitivities, varies within \mathcal{R} . For instance, $J(\hat{r})/\text{Vol}(\mathcal{R})$ for $\hat{r} = \theta_5$ at the maximum likelihood estimate $\theta_5 = -1.2015$ is 0.0011. This is about 400% and 80% smaller than the sensitivity for the maximal-margin and mean value respectively over \mathcal{R} . Hence, variation within \mathcal{R} caused by variations in ϵ will yield different sensitivity values and rankings. In the context of the maximum likelihood approach, a sensitivity metric to θ_i is $1/\bar{\sigma}_i$. This metric leads to values [46.78, 96.06, 48.13, 30.40, 37.66, 14.63] and an associated ranking of [2, 3, 1, 5, 4, 6] which differs from the two rankings above. The ratio between the most and the least important parameters is 6.56, 158.3, and 133 according to $1/\bar{\sigma}_i$, w_r , and μ_r respectively. These major differences are to be expected since the conventional approach is based on local sensitivities of the output y , whereas the global sensitivities proposed account for the effects of the model's parameters on several figures of merit over the entire validation domain.




Figure 9. $J(\hat{r})/\text{Vol}(\mathcal{R})$ for $\hat{r} = \theta_i$, for $i = 1, \dots, 6$.

VII. Discussion

In this section a qualitative comparison is made between conventional uncertainty quantification and model validation techniques and those proposed herein.

For the conventional approach of Section II, the dynamic model evaluated at the maximum-likelihood parameter estimate yields a prediction that optimally matches the measured data from a single flight test maneuver. As such, the identified parameter value is based on a single flight experiment and the particular performance metric in Eq. (5). There is no theory or rationale for assessing the error incurred by evaluating the dynamic model at any other realization of the model's parameters. The accuracy of the identified model parameters using this approach depends on the suitability of the model for representing the system dynamics, the characteristics of the measurement noise ν , and the effectiveness of the input in exciting the system response in the desired frequency range. The uncertainty set that results from applying this approach, which uses local sensitivities evaluated at the parameter estimate, has a statistical foundation solely based on the underlying assumptions prescribing the measurement noise. Data from multiple flight test maneuvers can be used to refine both the dynamic model and the uncertainty set, but this must be achieved by concatenating the data from the multiple maneuvers, or by using sequential calculations in a Bayesian formulation.¹

Conversely, the framework proposed here bounds the set of model parameters for which the empirical model is requirement compliant, e.g., the model prediction accurately describes the measured output corresponding to several maneuvers within an admissible level of error. The relationship between this set of parameter values (i.e., the validation domain) and the

prediction error is verifiable and objective. The validation domain can be systematically reduced by imposing additional requirements or by using data from additional flight test maneuvers. Both the structure of the dynamic model and the requirements can assume arbitrary forms, e.g., a nonlinear system model can be subject to a matching error based on an l_∞ -norm, or a linear system model can be subject to requirements in the frequency domain. The calculation of the validation domain is not based on local sensitivities and assumed probability density functions, but rather on the ability of the model to satisfy the requirements imposed upon it.

According to conventional methods, the value of $\bar{\sigma}_j$ is related to the inverse of the sensitivity of the model prediction $y(\theta)$ to variations in θ_j , i.e., high sensitivity to a particular parameter is reflected in a correspondingly tight uncertainty range, cf. Eq. (10). This sensitivity metric solely depends on a derivative evaluated at a single parameter realization, and thereby it is local in nature. In contrast, the sensitivities in Section IV evaluate the range of variability of the requirements functions as θ varies over the entire validation domain. The results of local and global sensitivity analyses will coincide when the dependence of g on θ is almost linear over the validation domain. Local and global sensitivities will likely differ in most other cases.

The variation in the model predictions due to uncertainty in the model parameters is commonly studied using Monte Carlo analysis. This analysis often uses a prescribed probabilistic uncertainty model for θ , typically in the form of a joint probability density function defined over some support set $\mathcal{S} \subset \mathbb{R}^n$. Ideally, the support set \mathcal{S} should not extend beyond the validation domain. This will guarantee that only the validated model parameters will be taken into account during the analysis. Unfortunately, the complex geometry of \mathcal{V} usually precludes this practice. The best alternative is for \mathcal{S} to contain the validation domain as tightly as possible. This can be attained by making \mathcal{S} equal to the outer bounding set of \mathcal{V} resulting from the free- m formulation proposed herein. While an analysis based on this uncertainty model considers all the validated dynamic models as desired, parameter values in \mathcal{S} that are not in \mathcal{V} , for which the model prediction is invalid, will also be included. When \mathcal{V} is not fully contained by \mathcal{S} , elements of the unknown parameter space that adequately predict the observations will be omitted from the study.

In the conventional approach the uncertainty set in Eqs. (12, 13) depends on the value η , which, in turn, depends on the analyst's choice for the confidence level. As such, the results and conjectures derived from a subsequent Monte Carlo analyses for which $\mathcal{S} = \Delta$ will be subjective. By contrast, analyses for which $\mathcal{S} = \tilde{\mathcal{B}}_{\text{outer}}$ will be exempt from this type subjectivity. To totally eliminate the subjectivity in the prescription of the support set, the center of $\tilde{\mathcal{B}}_{\text{outer}}$, θ_c , must be an additional design variable in Eq. (27). Note that the framework proposed does not relate the validation domain to the "true" value of θ , which could well be an unknown constant or an aleatory variable, nor does it yield information about likelihoods within \mathcal{V} . The probabilistic distribution within the support set \mathcal{S} , which is usually set according to past experience and expert judgement, will remain subjective.

Knowledge of the validation domain for empirical parametric models has many practical uses. For example, Monte Carlo analysis for evaluating mission success in the presence of modeling uncertainty requires this information. The accuracy of model uncertainty characterization affects not only the predicted probability of mission success, but also influences mission cost and risk. Because the unknown parameters in empirical dynamic models for aircraft typically represent aerodynamic characteristics (such as stability and control deriva-

tives), knowledge of the uncertainty in the parameter estimates is essential for evaluating correlation with results from ground-based aerodynamic prediction methods such as wind tunnel testing and computational fluid dynamics. The resulting comparisons can inform and guide efforts for improvement in the ground-based prediction methods by indicating which results are in agreement, based on model performance, and which exhibit a meaningful difference. Furthermore, many critical evaluations related to aircraft stability, control, and fault detection depend on deciding when the values of dynamic model parameters cross some threshold values. Distinguishing real parameter variations from random statistical variations associated with the parameter estimation process is critical for these tasks, and requires good uncertainty measures.

VIII. Concluding Remarks

This paper develops a framework for the parameter estimation and validation of empirical dynamic models subject to an arbitrary set of validation criteria. The validation requirements imposed upon the model, which might include several sets of input-output data and arbitrary specifications in time and frequency domains, are used to determine if model predictions are within admissible error limits. This analysis is used to calculate parameter estimates, to characterize the uncertainty in their values, and to validate the resulting empirical model against several figures of merit. In contrast to standard approaches in which the only source of discrepancy between the observations and the prediction is assumed to be measurement noise, this approach casts all discrepancies (e.g., those caused by measurement noise, model-form uncertainty, unsteady aerodynamics, etc.) as parametric uncertainty. This work sets forth a new paradigm for deriving, characterizing and validating empirical dynamic models.

References

- ¹Klein, V. and Morelli, E. A., *Aircraft system identification - theory and practice*, AIAA Education Series, Reston, VA, 1st ed., August 2006.
- ²Maine, R. E. and Iliff, K. W., “The theory and practice of estimating the accuracy of dynamic flight-determined coefficients,” NASA RP 1077, July 1981, pp. 1–31.
- ³Morelli, E. A. and Klein, V., “Determining the accuracy of maximum likelihood parameter estimates with colored residuals,” NASA CR 194893, March 1994, pp. 1–38.
- ⁴Morelli, E. A. and Klein, V., “Accuracy of aerodynamic model parameters estimated from flight test data,” *AIAA Journal of Guidance, Control and Dynamics*, Vol. 20, No. 1, 1997, pp. 74–80.
- ⁵Crespo, L. G., Giesy, D. P., and Kenny, S. P., “Robust analysis and robust design of uncertain systems,” *AIAA Journal*, Vol. 46, No. 2, 2008.
- ⁶Kenny, S. P., Crespo, L. G., and Giesy, D. P., “Dimensionality reduction for uncertain dynamic systems,” *International Journal of Numerical Methods in Engineering, Special issue on Uncertainty Quantification and Prediction Science*, Vol. 80, No. 1, March 2009, pp. 767–788.