International Workshop on Future of CFD and Aerospace Sciences

Gradient-based Methods for Uncertainty Quantification in Hypersonic Flows

Brian Lockwood[a], Dimitri Mavriplis[a,*]

[a] *University of Wyoming, Dept. 3295 1000 E. University, Laramie, WY and 82071, USA*

Abstract

In this paper, the use of gradient information for the acceleration of uncertainty quantification within the context of viscous hypersonic flows is examined. In particular, the variability of simulation outputs associated with uncertain parameters relating to physical models within the CFD simulation is predicted and gradient-based methods are used to reduce the cost of this prediction. The gradient of a simulation output is calculated via a discrete adjoint approach. By using an adjoint-based approach, the sensitivity of an objective to a large number of input parameters can be calculated in an efficient and timely manner. The additional information acquired from these derivative values are then leveraged to accelerate the quantification of the different forms of uncertainty: aleatory, epistemic, and mixed. For the aleatory case, gradient-enhanced surrogates, particularly Kriging models, are used to represent the dependence of the simulation output on input variables and is used as a basis for inexpensive Monte Carlo sampling. For epistemic uncertainties, a constrained, gradient-based optimization is used to determine the appropriate interval on the simulation output. Finally, for the mixed case, an optimization approach coupled with a surrogate method is used to generate the appropriate interval statistics. These strategies are demonstrated for a realistic CFD simulation utilizing a five species, two temperature real gas model with input parameters drawn from freestream conditions, transport relationships and the chemical kinetics model and the performance of these methods is assessed based on comparison with exhaustive sampling approaches.

Keywords: Uncertainty quantification, Adjoint Methods, Gradient-based Optimization, Hypersonic flow

1. Main text

Uncertainty quantification (UQ) is increasingly becoming an important part of computational science, providing the ability to assess the quality of computational results and apply confidence bounds to output metrics. This elevated importance has been driven in part on the increasing reliance of computational modeling within the design and analysis of complex engineering systems, particularly problems in which experimental data is difficult or impossible to obtain, such as the simulation of hypersonic flow. Various sources can contribute to the uncertainty within a simulation, such as measurement errors, modeling inadequacies[1], or manufacturing tolerances [2]. In the case of hypersonic flows, numerous constitutive relations are required, each of which with a number of experimentally derived constants and parameters. These parameters are often the result of experimental measurements and have an associated variability. In general, this variability can be either aleatory or epistemic or can have contributions from both. The traditional methods for quantifying both types of uncertainty require a large number of simulations, making traditional uncertainty quantification prohibitively expensive for complex simulations. Uncertainty quantification in its traditional form is even more intractable when it is included within the optimization process, where by each design iteration requires the uncertainty to be calculated. To address these limitations of UQ, numerous acceleration strategies have been proposed to reduce the number of simulations required for reliable statistics. In this paper, gradient-based methods for rapid uncertainty quantification are discussed for each type of uncertainty.

Aleatory uncertainties arise due to the inherent randomness of a variable and are characterized by a probability distribution [3]. For aleatory inputs, the goal of uncertainty quantification is to determine the distribution of an output quantity due to these input distributions. This quantification can be performed in a relatively straightforward, although expensive, manner. In previous work, Monte Carlo sampling has been used within the context of hypersonic flows to build up the required statistics for relevant simulation

*Corresponding author. Tel.: +1-307-766-2868 ; fax: +1-307-766-2695 .
Email address: mavripl@uwyo.edu . (Dimitri Mavriplis)

outputs [4, 5]. For this type of sampling, computing an output requires a complete computational fluid dynamic (CFD) simulation, making exhaustive sampling expensive for complex problems. When only a limited number of simulation outputs are of interest, a typical approach for reducing the expense of Monte Carlo sampling is the use of an inexpensive surrogate. This surrogate approximates the relationship between the true function value and the input parameters and is built based on a limited number of function evaluations. Because the surrogate is inexpensive to evaluate, exhaustive sampling can be performed to build the required statistics of the output. Surrogate models range in complexity from simple extrapolations [6, 7] to more sophisticated models, such as least-squares polynomials, multilayer perception, radial basis functions, and kriging. In computational fluid dynamics (CFD), kriging methods in particular have gained popularity [8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18] although their use within the context of uncertainty quantification for hypersonic flows is limited [19]. Techniques based on polynomial chaos have also been employed with success in the context of hypersonic flows [20, 21]. One drawback of surrogate based methods is the curse of dimensionality, whereby the number of samples required for an accurate surrogate increases exponentially as the number of input parameters grows. One method for overcoming this limitation is the incorporation of gradient information into the training of the surrogate [22, 12, 16, 17, 23, 24]. When adjoint methods are employed, this gradient can be evaluated with a cost approximately equal to the simulation of the physical problem [25, 26, 27]. By incorporating derivative values, the cost associated with training an accurate surrogate can be greatly reduced.

For most engineering simulation, the output uncertainty is typically the result of a relatively small number of variables. Because of this fact, dimension reduction strategies can also be used to extend the applicability of surrogate based approaches. One basis for this dimension reduction is sensitivity analysis. Two forms of sensitivity analysis are possible: localized analysis and global analysis (GSA). For a localized analysis, the magnitude of the derivative value is used to assess the importance of each variable. For hypersonic flow applications, this localized sensitivity analysis is demonstrated in Reference [28]. For problems where large input perturbations are possible (i.e. parameters with large uncertainties), this localized approach may no longer accurately predict the effect of these parameters on the output. Additionally, the localized approach cannot account for interference effects, where by the sensitivity of a given variable is altered by perturbation of the other variables. To account for these effects, global sensitivity analysis is performed [4]. This global sensitivity analysis is typically performed using Monte Carlo sampling and sensitivities are calculated based on the correlation between the input variables and the simulation output. Because of this reliance on Monte Carlo, the cost associated with this global sensitivity analysis is prohibitively expensive for complex simulations. As was the case for aleatory uncertainty, a surrogate can be built and used as a basis for inexpensive Monte Carlo sampling; however, this surrogate must extend to large dimension without a dramatic increase in cost. One such surrogate is polynomial regression enhanced with derivative values [23]. For this work, a low-order version of this surrogate provides the basis for a rapid global sensitivity analysis which in turn provides the necessary means of dimension reduction required for the aleatory uncertainty quantification.

Epistemic uncertainty arises from a lack of knowledge regarding the true value of a parameter and is typically specified by using an interval. The goal of uncertainty quantification for epistemic uncertainties is to determine the interval output of a quantity due to specified input intervals. The quantification of epistemic uncertainties has been scarcely explored in the context of hypersonic flows. This situation is in spite of the fact that epistemic uncertainties are the dominant form of uncertainty present in hypersonic flows and previous studies assuming pure aleatory uncertainties, although important initial steps, have likely underestimated the uncertainty associated with simulation objectives [4, 29]. Epistemic uncertainty may be quantified via sampling based approaches or via optimization. Typically, Latin hypercube sampling is used for epistemic uncertainties, although other methods such as approaches based on random sampling and Dempster-Shafer evidence theory can be used [30, 31, 32]. For Latin hypercube sampling in particular, the required number of samples grows quickly as the dimension of the problem increases, making the quantification of epistemic uncertainties for large-dimension problems difficult [3]. As was the case with aleatory uncertainty, one possible solution is to replace sampling with a surrogate model; however, this approach will again eventually encounter the curse of dimensionality as the input dimension increases. The other main approach for epistemic uncertainty quantification is to pose the problem as a constrained

optimization problem. That is, given input parameters within specified ranges, determine the maximum and minimum values of an output function. Although this approach entails solving a complicated global optimization problem with the possibility of multiple extrema, the number of function evaluations to solve the optimization problem scales more readily to high-dimensional problems if a gradient-based optimizer is employed [33].

The problem of epistemic uncertainty quantification is further complicated when contributions from aleatory sources are also considered. This mixed aleatory/epistemic uncertainty quantification typically relies on a nested sampling strategy. Although the required number of samples grows extremely fast, these strategies are conceptually easy to understand and are capable of separating the effects of each type of uncertainty [34, 3]. For nested strategies, samples are first drawn from the epistemic variables; and for each set of epistemic variables, the distribution of the output due to the aleatory variables is determined based on sampling over the aleatory variables. Since the number of samples required for the epistemic uncertainty grows exponentially fast, the expense of nested sampling grows rapidly with respect to the number of epistemic variables[3]. For hypersonic flows, the number of epistemic variables is typically much greater than the number of aleatory variables. Hence, for complex models with many uncertain epistemic variables, nested approaches will quickly become prohibitively expensive. Here, too, surrogates can be created as a function of all variables and samples extracted according to a nested strategy. For relatively low dimensions, this strategy can be effective [21] and, when combined with gradient-enhancement, could be applied to problems of moderate dimension. However, once the number of epistemic variables increases sufficiently, surrogate-based approaches will again become prohibitively expensive as the required number of samples increases for an accurate surrogate. In order to address this concern, combination sampling/optimization approaches have been explored[34]. For mixed aleatory/epistemic problems, the goal of the uncertainty quantification is to produce a region in which the function is contained with a specific level of confidence, known as a P-Box [3]. Stated in other terms, the bounds of the confidence interval of the output distribution must themselves be an interval in order to account for the epistemic uncertainties. Because the details of this bounding box are irrelevant and only the bounds of this box are required, the sampling with respect to the epistemic variables may be replaced by optimization. In principle, these mixed optimization/sampling approaches may be posed in two ways: determining intervals of statistics and determining statistics of intervals.

The first approach can be viewed as an optimization-under-uncertainty problem with the metric of the optimization defined as a relevant statistic of the aleatory distribution, such as the mean and variance, bounds on a confidence interval, or a reliability index. For each step in the optimization, the aleatory uncertainty is quantified, typically by means of a polynomial chaos surrogate model, and the relevant statistic of the distribution is calculated and used as the metric for the optimization [34]. For the second approach (referred to as statistics-of-intervals, or SOI), an optimization problem is posed for each set of aleatory variables, and repeated optimizations are used to determine the relevant statistics of the interval. To reduce the number of required samples, a surrogate model of the optimization results is constructed with respect to the aleatory variables, ensuring few total optimizations are required to characterize the statistics of the interval. For either of these methods, a gradient-based optimization can be used, reducing the cost of each optimization and ensuring optimal scaling as the input dimension increases. The pros and cons of each method as applied to hypersonic flows are discussed in Reference [35]. In this paper, the statistics-of-intervals approach is discussed exclusively.

The structure of this paper is as follows. First, the flow solver used throughout this work is explained and demonstrated for the test case used throughout this paper. Second, because gradient-based methods are the focal point of this work, the derivation of the adjoint sensitivity equations is given. To provide a basis for the dimension reduction required for the uncertainty quantification strategies, the use of this derivative information within a global sensitivity analysis is then outlined and example results presented. Finally, with these derivative values and sensitivity analysis established, gradient-based techniques for aleatory, epistemic and mixed uncertainties are explained and demonstrated for a real gas simulation.

1.1. Flow Solver

In this section, the physical models and CFD flow solver used in this paper are outlined. For this work, the Navier-Stokes equations are solved numerically in two dimensions via a cell-centered finite-volume

scheme on unstructured meshes using triangular and/or quadrilateral elements. In vector form, the Navier-Stokes equations are given by

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \vec{F}(\mathbf{U}) = \nabla \cdot \vec{F}_v(\mathbf{U}) + \mathbf{S}(\mathbf{U}), \tag{1}$$

where $\mathbf{U}$ are the conserved flow variables, $\vec{F}$ is the inviscid flux, $\vec{F}_v$ is the viscous flux, and S contains any source terms required for the physical model, such as reaction or energy coupling terms. For this work, a nonequilibrium real gas physical model is examined. The real gas model is a five-species, two-temperature model for non-ionizing air [36]. Both the Dunn-Kang chemical kinetics model and the Park model have been implemented for this work. For use within uncertainty quantification, the Dunn-Kang model is used due to the ease with which uncertainty parameters are specified, while the Park model is used for the purposes of solver verification. For the Dunn-Kang model, the forward and backward reaction rates are specified directly by means of an Arrhenious relation, given as:

$$K_f = C_f T_a^{\eta_f} e^{-\frac{E_{a,f}}{k_B T_a}} \tag{2}$$

$$K_b = C_b T_a^{\eta_b} e^{-\frac{E_{a,b}}{k_B T_a}} \tag{3}$$

where $E_{a,f}$ and $E_{a,b}$ represent the activation energy for the forward and backward reactions respectively, k_B is Boltzmann's constant, and T_a is a characteristic temperature for the reaction. The specific heats are calculated via fourth order polynomial curve fits covering various temperature ranges. The total enthalpy is calculated simply by integrating these curve fits and incorporating the proper heat of formation information [37]. The transport model is a collision integral model. For this model, viscosity, thermal conductivity, and diffusion coefficients are calculated based on linear interpolation of collision integrals between 2000 K and 4000 K[37, 38].

$$log_{10}(\Omega_{s,r}^{k,k}) = log_{10}(\Omega_{s,r}^{k,k})_{2000} + \left[log_{10}(\Omega_{s,r}^{k,k})_{4000} - log_{10}(\Omega_{s,r}^{k,k})_{2000}\right] \frac{ln(T) - ln(2000)}{ln(4000) - ln(2000)} \tag{4}$$

Here, $\Omega_{s,r}^{k,k}$ represents the collision integral between species s and species r. For the five-species model, 15 independent collision interactions are possible, giving a total of 60 parameters, since two separate collision integrals ($\Omega_{s,r}^{1,1}$ and $\Omega_{s,r}^{2,2}$) are used at each temperature. The complete real gas model used in this work contains approximately 250 parameters, embedded within the constitutive models for the reaction rates, transport coefficients, relaxation times, and caloric equations of state.

In order to solve problems using the above model, a two-dimensional, cell-centered finite-volume code was written. The governing equations described above are first discretized in space, and the solution is advanced in time to steady state using a fully implicit approach. In semi-discrete form, the equations have the following form:

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{R}(\mathbf{U}) = 0 \tag{5}$$

The residual within each cell is given by the sum of the normal inviscid and viscous fluxes over all faces plus a cell centered contribution due to source terms. The inviscid flux is calculated by using gradient reconstruction of primitive variables, and gradients are calculated using Green-Gauss contour integration over the cell. The limiter used in this code is a combination of a pressure switch and the smooth Van Albada limiter[36, 38, 39]. The AUSM+UP flux function is used because of the ease with which it can be extended to additional equations and its applicability across a wide range of Mach numbers. In order to extend this flux function to the real gas model, a frozen speed of sound is used [39, 40].

The result of the spatial discretization outlined above is a system of coupled ODE's that are solved implicitly using a first order backward difference discretization. The result of this temporal discretization is a system of nonlinear equations that are solved by using an inexact Newton's method. This inexact method employs a number of approximations to improve the performance of the nonlinear solver. Instead

of solving the nonlinear system exactly, a fixed number of Newton iterations are performed. Additionally, instead of inverting the exact Jacobian, an approximate first-order Jacobian is used and is inverted iteratively. For this work, the Jacobian corresponding to the van Leer-Hanel flux function is used, and the Jacobian is inverted by using either a point or a line implicit approach. Once startup transients have been overcome, the preconditioner and transport quantities can be frozen during the pseudo-time step. A line-preconditioned GMRES solver is then used in an exact Newton method to accelerate the convergence of the solver to machine zero.

Table 1. Benchmark flow conditions

$V_\infty =$	5 km/s
$\rho_\infty =$	0.001 kg/m^3
$T_\infty =$	200 K
$T_{wall} =$	500 K
$M_\infty =$	17.605
$Re_\infty =$	376,930
$Pr_\infty =$	0.72

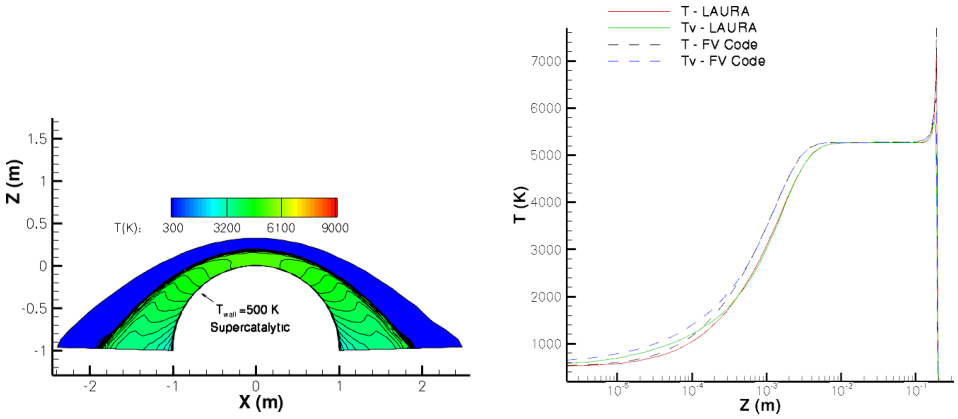

Figure 1. Validation of solver for 5 km/s flow over circular cylinder. Left: Computed flow field temperature contours. Right: Comparison of temperatures along centerline with LAURA [41] results running on equivalent mesh.

The solver described previously was validated using the standard test case of 5 km/s flow over a circular cylinder with a super-catalytic, fixed-temperature wall. The conditions for this test case can be found in Table 1. The results of this test case were compared with those of the well-validated code LAURA[41, 42] and are depicted in Figures 1 and 2. For these comparisons, the Park chemical kinetics model was used. Although this model shows better agreement with the validation codes, the Dunn-Kang model is ultimately used for all of the demonstration sensitivity and uncertainty quantification results. As the figure shows, the solver is able to match the LAURA validation results closely.

1.2. Adjoint Derivation

Presented next is an outline of the sensitivity procedure used to calculate the gradient of an objective. Further details of this adjoint implementation are presented elsewhere [28]. To determine the gradient, the code is differentiated piece by piece, and the final sensitivity is constructed by using the chain rule. In order to illustrate this process, the following objective functional dependence is considered.

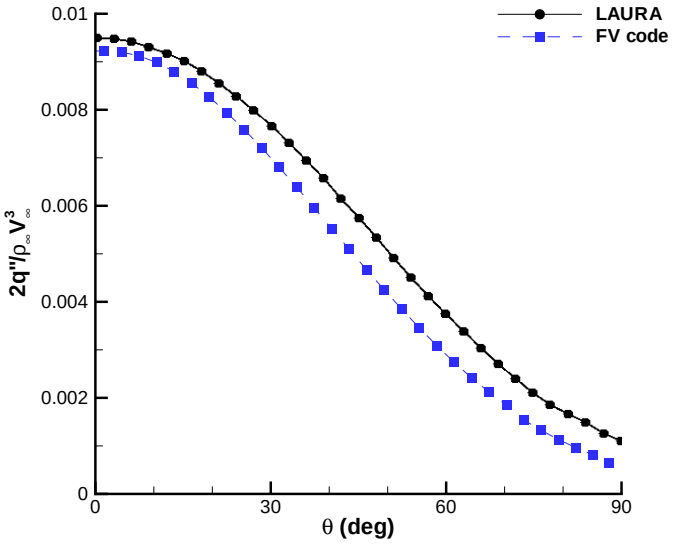

Figure 2. Validation of solver for 5 km/s flow over circular cylinder: Surface heating distribution compared to the LAURA result

$$L = L(D, \mathbf{U}(D)) \tag{6}$$

In addition to this objective, a constraint is needed. For the steady problems considered in this work, the constraint is that the spatial residual must equal zero.

$$\mathbf{R}(D, \mathbf{U}(D)) = 0 \tag{7}$$

Both the constraint and the residual have an explicit dependence on the input parameters, or design variables D, and an implicit dependence through the flow variables $\mathbf{U}$. In order to determine the sensitivity derivative, the objective can be differentiated by using the chain rule as follows. [43]

$$\frac{dL}{dD} = \frac{\partial L}{\partial D} + \frac{\partial L}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial D} \tag{8}$$

The constraint may be differentiated in a similar manner. In this case, the derivative is equal to zero, since the constraint must be satisfied for all admissible values of D and $\mathbf{U}$.

$$\frac{\partial \mathbf{R}}{\partial D} + \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial D} = 0 \tag{9}$$

Solving for $\frac{\partial \mathbf{U}}{\partial D}$, and substituting into the objective derivative, one obtains the forward sensitivity equation (also known as the tangent linear model).

$$\frac{dL}{dD} = \frac{\partial L}{\partial D} - \frac{\partial L}{\partial \mathbf{U}} \frac{\partial \mathbf{R}}{\partial \mathbf{U}}^{-1} \frac{\partial \mathbf{R}}{\partial D} \tag{10}$$

As the equation shows, the residual Jacobian must be inverted once for each design variable. However, the same $\frac{\partial \mathbf{U}}{\partial D}$ may be used for each objective L. Because of the expense associated with inverting the residual Jacobian, the forward sensitivity approach is best suited for problems with few design variables and multiple objectives.

The adjoint sensitivity equation is found by taking the transpose of the forward equation

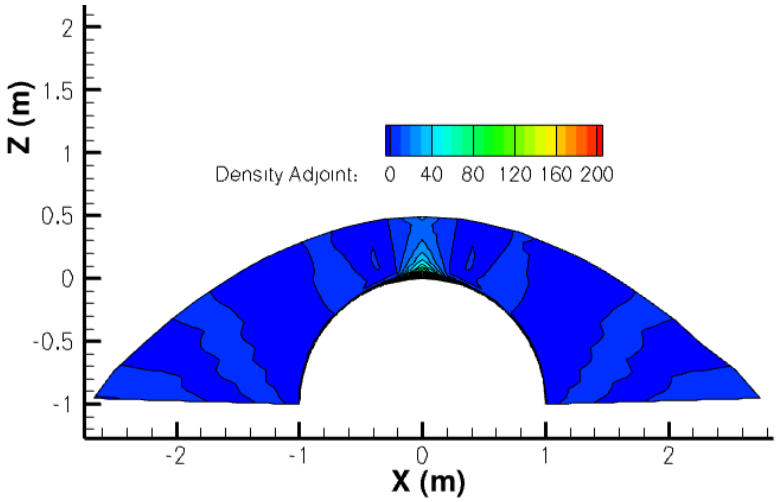

Figure 3. Density adjoint for integrated surface heating

$$\frac{dL}{dD}^T = \frac{\partial L}{\partial D}^T - \frac{\partial \mathbf{R}}{\partial D}^T \frac{\partial \mathbf{R}}{\partial \mathbf{U}}^{-T} \frac{\partial L}{\partial \mathbf{U}}^T \tag{11}$$

where the last two terms can be replaced by the adjoint variable $\mathbf{\Lambda}$, defined as follows.

$$\frac{\partial \mathbf{R}}{\partial \mathbf{U}}^T \mathbf{\Lambda} = -\frac{\partial L}{\partial \mathbf{U}}^T \tag{12}$$

A sample adjoint solution for the 5 km/s benchmark is found in Figure 3. This figure shows the adjoint variable for surface heating associated with the density. The magnitude of this variable roughly represents the importance of the flow variable on the objective of interest. As expected for surface heating, the adjoint variable is largest near the surface of the cylinder. With this definition of the flow adjoint, the final sensitivity equation is given by the following.

$$\frac{dL}{dD} = \frac{\partial L}{\partial D} + \mathbf{\Lambda}^T \frac{\partial \mathbf{R}}{\partial D} \tag{13}$$

Determining the solution of the flow adjoint equation roughly follows the procedure used to solve the analysis problem. A simplified preconditioner matrix is used to advance the adjoint solution in a defect-correction scheme [43]. The effect of the exact Jacobian required for the defect-correction solver is built up by using automatically differentiated subroutines. The automatic differentiation used in this work is provided by the Tapenade Automatic Differentiation Engine [44]. Using these automatically differentiated subroutines, a line-preconditioned GMRES solver is used to invert the transpose of the Jacobian. With this adjoint implementation, the sensitivity of an objective to any number of parameters can be computed with a nearly constant amount of work.

1.3. Sensitivity Analysis

When calculated using an adjoint technique, the gradient provides additional information about the simulation with limited additional overhead. One application for this additional information is sensitivity analysis, where the effect of each input parameter on the output can be quantified and the most important parameters on the output can be determined. In addition to being valuable in its on right, sensitivity analysis is often critical for rapid uncertainty quantification techniques as the expense of these techniques often

increases dramatically as the dimension of the problem increases. Hence, for these methods, sensitivity analysis represents a means for dimension reduction. While a localized sensitivity analysis for this application was presented in Reference [28], this work will utilize a global sensitivity analysis.

For a global sensitivity analysis, Monte Carlo sampling is performed and correlation coefficients between the input and output are calculated. Although the work associated with this sampling is independent of the number of design variables, the expense of this sampling is prohibitively high for complex simulations due to the slow convergence of output statistics. The correlation coefficient for variable D_i is given by the following [4]:

$$r_i = \frac{cov(D_i, y)}{\sigma_{D_i}\sigma_y} \tag{14}$$

Here, y represents the output of interest from the simulation, σ_{D_i} represents the standard deviation of the input design variable and σ_y represents the standard deviation of the output. The standard deviation of the input design variable is a quantity that must be taken from the relevant literature or estimated based on some expert judgement or experience [3, 4]. The quantity σ_y is measured empirically from the Monte Carlo data set. For the purposes of ranking the sensitivities, the square of the correlation coefficient is used. In addition to acting as a proxy for the magnitude of the sensitivity, the square of the correlation coefficient represents the fraction of the output variance from each of the input parameters [4].

In principle, the process of global sensitivity analysis can be accelerated through the use an inexpensive surrogate that can be used to approximate the design space and Monte Carlo sampling can be performed on this surrogate. However, for most surrogate models, the expense associated with constructing an accurate surrogate increases exponentially as the dimension of the design space increases. This fact makes the use of traditional surrogate models for global sensitivity analysis ineffective for large dimensional problems. Although traditional surrogate models may not be useful for GSA, the incorporation of gradient information into the surrogate model can greatly reduce the cost associated with training an accurate surrogate and mitigate the cost increases as the dimension rises [23, 22, 24].

For this work, a simple polynomial regression model enhanced with derivative values is used as the basis for a rapid global sensitivity analysis [23]. For a regression-based model, the output is a linear combination of polynomials of varying orders, depicted in Equation 15.

$$y(D) = \sum_s \beta_s \Psi_s(D) \tag{15}$$

Here, $\Psi(D)$ represent a series of polynomials in D with degree less than some specified order p and β are a set of undetermined coefficients. For this work, hermite polynomials are used as the basis and the multidimensional polynomial is constructed by means of a tensor product [23, 45]. Based on this basis and the simulation results, the weights in this linear combination are determined via a least-squares process. Hence, the amount of information required to construct the model is proportional to the number of terms in the basis. For a general polynomial order and dimension, the size of the basis is given by [45]:

$$S = \frac{(d+p)!}{d!p!} \tag{16}$$

where d is the dimension of the space and p is the maximum polynomial order. In practice, the amount of information used to construct the regression is greater than the size of the basis, typically by a factor of two, and the coefficients in the regression are solved for using least-squares. For a regression based on function values alone, a new simulation is required for each piece of information, making the cost of constructing this model increase rapidly with dimension and polynomial order. However, when gradient information is available, each simulation result produces $d + 1$ pieces of information with a cost approximately equal to the original simulation. When gradient information is included in the regression training, the number of simulations required to create the regression is given as:

$$N \geq \lceil \frac{(d+p)!}{d!p!(d+1)} \rceil \tag{17}$$

Table 2. Input Model Parameters with Uncertainty

Number	Variable	Mean	Standard Deviations
1	$\rho_\infty\ (kg/m^3)$	1×10^{-3}	5%
2	$V_\infty(m/s)$	5000	15.42
3-17	$A^{1,1}_{s-r}$	1	5%
18-32	$A^{2,2}_{s-r}$	1	5%
33-49	ξ_f	0	0.25
50-66	ξ_b	0	0.25

For this work, the polynomial order is limited to second order ($p = 2$), giving only a linear increase in the cost associated with training the regression as the dimension increases.

$$N \geq \frac{(d+2)}{2} \tag{18}$$

Incorporating derivative observations into the creation of the regression is relatively straight-forward, requiring only the differentiation of Equation 15 and incorporating these equations into the collocation matrix inverted to determine the regression coefficients, illustrated in Equation 19 [23].

$$\begin{bmatrix} \Psi_1(D_1) & \Psi_2(D_1) & \cdots & \Psi_s(D_1) \\ \frac{\partial \Psi_1(D_1)}{\partial D_1} & \frac{\partial \Psi_2(D_1)}{\partial D_1} & \cdots & \frac{\partial \Psi_s(D_1)}{\partial D_1} \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial \Psi_1(D_1)}{\partial D_d} & \frac{\partial \Psi_2(D_1)}{\partial D_d} & \cdots & \frac{\partial \Psi_s(D_1)}{\partial D_d} \\ \vdots & \ddots & \ddots & \vdots \\ \Psi_1(D_N) & \Psi_2(D_N) & \cdots & \Psi_s(D_N) \\ \frac{\partial \Psi_1(D_N)}{\partial D_1} & \frac{\partial \Psi_2(D_N)}{\partial D_1} & \cdots & \frac{\partial \Psi_s(D_N)}{\partial D_1} \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial \Psi_1(D_N)}{\partial D_d} & \frac{\partial \Psi_2(D_N)}{\partial D_d} & \cdots & \frac{\partial \Psi_s(D_N)}{\partial D_d} \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_s \end{bmatrix} = \begin{bmatrix} y(D_1) \\ \frac{\partial y(D_1)}{\partial D_1} \\ \vdots \\ \frac{\partial y(D_1)}{\partial D_d} \\ \vdots \\ y(D_N) \\ \frac{\partial y(D_N)}{\partial D_1} \\ \vdots \\ \frac{\partial y(D_N)}{\partial D_d} \end{bmatrix} \tag{19}$$

Using this gradient-enhanced polynomial regression, a GSA based on a second order regression can be constructed with a cost that increases only linearly as the dimension increases. Hence, this method can be applied even for large dimension problems. To demonstrate this global sensitivity analysis and provide a basis for further dimension reduction within the uncertainty quantification, a regression-based global sensitivity analysis was performed on the $5km/s$ benchmark problem outlined in Section 1.1. For this sensitivity analysis, 66 of the input parameters to the simulation are examined, drawn from freestream conditions, transport models and reaction rate specification. The effect these input parameters have on the integrated surface heating is quantified. For the second order regression, 68 function/gradient evaluations are used to train the model, computed from input samples distributed uniformly throughout the design space via Latin Hypercube sampling. For this analysis, these parameters are assumed to follow a Gaussian distribution. The results of the global sensitivity analysis are compared against Monte Carlo sampling using 6331 simulation results based on the square of the correlation coefficient. The variables examined for this sensitivity analysis are listed in Table 2.

For the transport coefficients, the collision integrals are treated as the parameters of interest and the variables examined for this analysis take the form of a multiplicative constant on the input collision integrals $\hat{\Omega}^{k,k}_{s,r}$ [5].

$$\Omega^{k,k}_{s,r} = A^k_{s,r} \hat{\Omega}^{k,k}_{s,r} \tag{20}$$

Table 3. Top 10 Parameters from P=2 Regression Sensitivity Analysis compared with Global

Rank	Variable	Monte Carlo Rank	Regression Contribution	Monte Carlo Contribution
1	ρ_∞	1	0.56879	0.60055
2	$O_2 + O \leftrightarrows 2O + O$ (f)	2	1.0002×10^{-1}	1.0610×10^{-1}
3	$O_2 + O_2 \leftrightarrows 2O + O_2$ (b)	6	5.7669×10^{-2}	2.1621×10^{-2}
4	$NO + O \leftrightarrows N + O + O$ (b)	3	4.0057×10^{-1}	5.1914×10^{-2}
5	N2-N2 (k=1)	5	3.7461×10^{-2}	3.1617×10^{-2}
6	O2-N2 (k=1)	4	3.3299×10^{-2}	4.2121×10^{-2}
7	N2-N2 (k=2)	8	2.1163×10^{-2}	1.9019×10^{-2}
8	O-N2 (k=2)	9	1.7395×10^{-2}	1.3874×10^{-2}
9	V_∞	14	1.3497×10^{-2}	4.8401×10^{-3}
10	$O_2 + O \leftrightarrows 2O + O$ (b)	13	1.1734×10^{-2}	7.4280×10^{-3}

With the variable defined in this manner, the sensitivity of 30 total parameters relating to the transport model is assessed.

For the reaction rates, the variable of interest for the sensitivity analysis is again a multiplicative constant on the reaction rates. Because of the large uncertainties typical of reaction rates, the variable represents the exponential on the multiplicative constant [4], given as:

$$K_f = 10^{\xi_f} C_{f,o} T_a^{\eta_f} e^{-\frac{E_{a,f}}{k_B T_a}} \tag{21}$$

$$K_b = 10^{\xi_b} C_{b,o} T_a^{\eta_b} e^{-\frac{E_{a,b}}{k_B T_a}} \tag{22}$$

Here, $C_{f,o}$ and $C_{b,o}$ represent the unperturbed coefficients, and ξ_f and ξ_b represent the parameters of interest for the forward and backward rates. This parametrization gives a total of 34 reaction rate variables.

The results of this regression-based GSA are summarized in Table 3 and compared to the corresponding result from the Monte Carlo analysis. In this Table, the ten most influential parameters on the surface heating are identified based on the regression-based GSA and compared with the ranking from the Monte Carlo analysis. Additionally, the square of the correlation coefficient predicted by each method is compared.

Based on the results in Table 3, two conclusions can be made. First, the regression-based analysis produces parameter rankings and uncertainty contributions in relatively good agreement with the Monte Carlo results at a fraction of the cost (68 function/gradient results vs. 6331 function results). Second, both the Monte Carlo and regression-based GSA indicate that the majority of the output uncertainty, measured by the square of the correlation coefficient, is the result of a small number of parameters, with these top ten accounting for over 90% of the output variance. Because the output variance is the result of a handful of variables, this sensitivity analysis can provide the justification for the dimension reduction used within the uncertainty quantification discussed in this work.

1.4. Gradient-based Uncertainty Quantification

In this section, gradient-based strategies for each of the different types of uncertainty are presented and demonstrated for the $5km/s$ real gas test case presented in Section 1.1. For the case of aleatory uncertainty, a gradient-enhanced Kriging model, augmented with a regression mean function, is used as a basis for inexpensive Monte Carlo sampling. For the case of epistemic uncertainty, constrained gradient-based optimization is used to determine the appropriate output interval based on the input intervals. Finally, for the mixed case, the statistics-of-intervals approach uses a combination of optimization and Kriging to determine intervals for statistics of interest.

Table 4. Variables for the Kriging Model based on Regression Global Sensitivity Analysis

Rank	Variable	Uncertainty Contribution	Total Contribution
1	ρ_∞	0.56879	0.56879
2	$O_2 + O \leftrightarrows 2O + O$ (f)	1.0002×10^{-1}	0.66882
3	$O_2 + O_2 \leftrightarrows 2O + O_2$ (b)	5.7669×10^{-2}	0.72649
4	$NO + O \leftrightarrows N + O + O$ (b)	4.0057×10^{-1}	0.76654
5	N2-N2 (k=1)	3.7461×10^{-2}	0.80400
6	O2-N2 (k=1)	3.3299×10^{-2}	0.83730
7	N2-N2 (k=2)	2.1163×10^{-2}	0.85847
8	O-N2 (k=2)	1.7395×10^{-2}	0.87586
9	V_∞	1.3497×10^{-2}	0.88936
10	$O_2 + O \leftrightarrows 2O + O$ (b)	1.1734×10^{-2}	0.90109
11	NO-N2 (k=1)	9.0220×10^{-3}	0.91011
12	$O_2 + N_2 \leftrightarrows 2O + N_2$ (f)	5.5034×10^{-3}	0.91562
13	$NO + O \leftrightarrows N + O + O$ (f)	5.1194×10^{-3}	0.92074
14	$NO + N_2 \leftrightarrows N + O + N_2$ (b)	2.9103×10^{-3}	0.92365
15	$NO + O \leftrightarrows O_2 + N$	2.6108×10^{-3}	0.92626
16	$NO + N \leftrightarrows N + O + N$ (b)	2.0229×10^{-3}	0.92828
17	$NO + N \leftrightarrows N + O + N$ (f)	1.9061×10^{-3}	0.93019

1.4.1. Gradient-Enhanced Kriging for Aleatory Uncertainty

To quantify the uncertainty associated with aleatory variables, a gradient-enhanced Kriging surrogate model is employed and used as a basis for inexpensive Monte Carlo sampling. In particular, the uncertainty of surface heating based on the 66 input parameters outlined in Section 1.3 was calculated, assuming each variable obeys a Gaussian distribution. To address the issue of extending the surrogate to large dimension, a combination of regression and Kriging is used. The regression model described in the previous section is first used to model the mean behavior of the Kriging model. Then, based on the results of the global sensitivity analysis, a Kriging model is built over the variables contributing the most to the output variance. This Kriging model is then used to enhance the predictions of the regression model. This functional dependence is represented as:

$$y(x) = m(x) + N(0, K(\tilde{x}, \tilde{x}')) \tag{23}$$

where $m(x)$ is the regression mean function and $N(0, K(\tilde{x}, \tilde{x}'))$ is a zero mean Gaussian process built over a subset of the variables. This Gaussian process was built over a set of variables accounting for 93% of the output variance (the target of 95% was not possible for the regression-based GSA so it was altered downward). This set consisted over the 17 most important variables and is given in Table 4.

The construction of the Kriging model is based on the difference between the training data and the mean function at the training location. For the gradient-enhanced model, this difference is also applied to the derivative values. Hence, the Gaussian process assumption applies to the following variables:

$$\begin{bmatrix} R(x) \\ \nabla R(x) \end{bmatrix} = \begin{bmatrix} y(x) - m(x) \\ \nabla y(x) - \nabla m(x) \end{bmatrix} = N(0, \underline{K}(x, x')) \tag{24}$$

For this work, the mean function of the Gaussian process is assumed to be zero, corresponding to simple kriging. The parameters required in the covariance matrix are determined by maximizing the marginal likelihood equation for the model. Throughout this work, the Matern class of covariance functions was used with the smoothness parameter ν set to $3/2$ since this combination of function and smoothness parameter produced the most accurate results[46, 47]. The output of a kriging model is also a Gaussian process and has an associated mean and variance. With the observation data, represented by the vector $\mathbf{y}$, the mean behavior

Table 5. Statistics predictions based on Kriging model compared with regression and Monte Carlo

	Regression	Kriging/Regression	Monte Carlo
Mean	1.0370E-002	1.0384E-002	1.0401E-002
Variance	8.6692E-008	9.2394E-008	9.7517E-008
Standard Deviation	2.9444E-004	3.0396E-004	3.1228E-004
95% Confidence Interval	5.6786%	5.8543%	6.0049%

of y away from these observations is predicted for a simple kriging model using an explicit mean function using the following equation.

$$y^* = m(D^*) + k(D^*, D)K^{-1}\left(\mathbf{y} - m(D^*)\right) \tag{25}$$

Here, $k(D^*, D)$ is the covariance between the test point of interest D^* and the training data D. The construction of Kriging models is detailed extensively in Reference [47] and the incorporation of gradient information into Kriging models is detailed in Reference [46, 48]. The performance of the Kriging/regression model is assessed based on a validation data set produced by Monte Carlo sampling. For this sampling, 6331 simulations were performed and statistics were calculated based on this data. The training of the Kriging model was based on 68 function/gradient evaluations sampled uniformly from the design space via Latin Hypercube sampling, the same training points used for the GSA.

Table 5 contains the statistic predictions based on the Kriging model and compared to the Monte Carlo results. In addition to presenting results for the Kriging based approach, statistic predictions based on the $P = 2$ regression alone are given to demonstrate the effect of enhancing the regression with the Kriging model. As the Table shows, although the regression model alone produces statistics in reasonable agreement with the Monte Carlo results, the Kriging model is able to provide improved predictions with no addition simulations, increasing the accuracy of the variance prediction. In addition to predicting distribution statistics, the CDF curve and associated quantiles predictions produced by the Kriging model is compared to that of Monte Carlo sampling. The CDF curves for the Kriging model and Monte Carlo are plotted in Figure 1.4.1. In addition to comparing the CDF curves directly, a QQ plot of the Kriging results vs. the Monte Carlo results is given in Figure 1.4.1. This plot compares the quantile predictions for the two methods, with a line of slope 1 indicating exact agreement between the two results.

As the plots demonstrate, the distribution predictions from the Kriging model match those from Monte Carlo extremely well. Further demonstrates of Kriging models for use in aleatory uncertainty quantification can be found in References [19, 24, 49].

1.4.2. Gradient-based Optimization for Epistemic Uncertainty

For the case of epistemic uncertainty, gradient-based optimization is used to produce the corresponding output interval by performing a minimization and maximization. The interval produced by optimization was validated by performing Latin hypercube sampling over the epistemic variables. Given the expense associated with validating the epistemic results, the dimension of the problem was reduced from 66 to 8 dimensions. Additionally, the uncertain variables were limited to the most sensitivity transport parameters as varying these parameters reliability lead to successful simulations even for large perturbations. The transport parameters used for this uncertainty study were chosen based on the sensitivity analysis in Section 1.3. The variables used for the epistemic uncertainty quantification are given in Table 6. The uncertainty for each variables was taken from reference [5] and the optimization itself was performed using bound constrained L-BFGS [50].

Figure 6 shows the convergence of the two optimization problems for the real gas solver. In addition to the optimization convergence, the bounds produced by Latin hypercube sampling are also plotted for validation purposes. The Latin hypercube sampling uses three samples in each direction, requiring 3^8 or 6561 simulation results. The optimization results provide conservative estimates for the bounds established

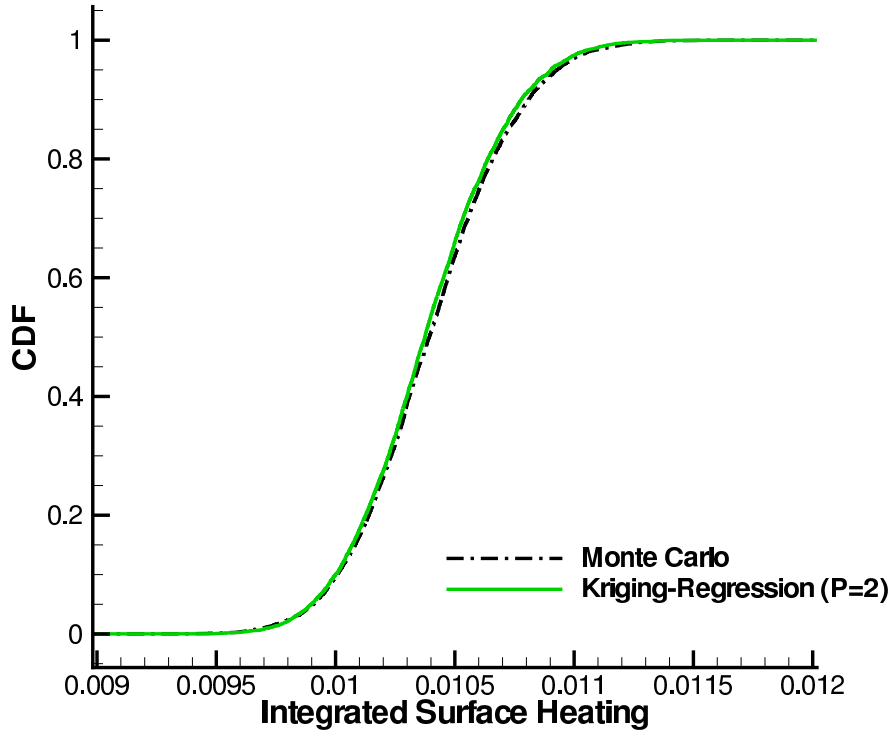

Figure 4. CDF prediction for Kriging/Regression model compared with Monte Carlo result.

Table 6. Epistemic Model Parameters

Variable	Unperturbed Value	Lower Bound	Upper Bound
$A^{1,1}_{N2-N2}, A^{2,2}_{N2-N2}$	1	0.8	1.2
$A^{1,1}_{N2-N}, A^{2,2}_{N2-N}$	1	0.8	1.2
$A^{1,1}_{N2-O}, A^{2,2}_{N2-O}$	1	0.8	1.2
$A^{1,1}_{N2-O2}, A^{2,2}_{N2-O2}$	1	0.8	1.2

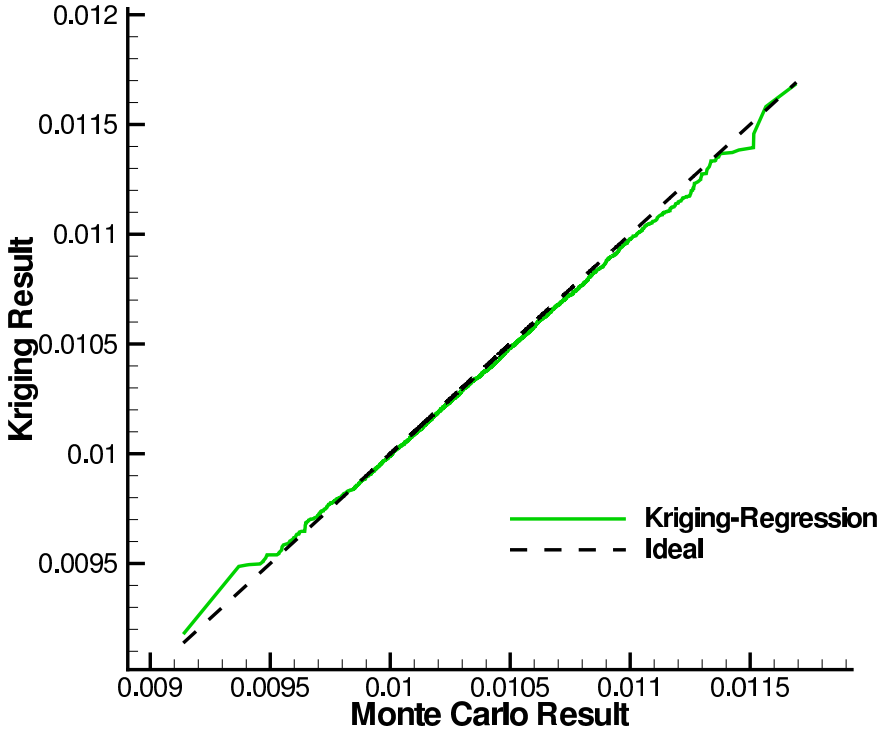

Figure 5. QQ plot of Kriging/Regression quantile predictions vs. Monte Carlo results. Exact agreement with Monte Carlo results indicated by line of slope 1.

Table 7. Uncertain Model Parameters

Variable	Type	Uncertainty
$\rho_\infty\ (kg/m^3)$	Aleatory	$\pm 10\%\ (\sigma = 5\%)$
$V_\infty (m/s)$	Aleatory	$\pm 30.84\ (\sigma = 15.42)$
$\Omega^{1,1}_{N2-N2}, \Omega^{2,2}_{N2-N2}$	Epistemic	$\pm 20\%$
$\Omega^{1,1}_{N2-N}, \Omega^{2,2}_{N2-N}$	Epistemic	$\pm 20\%$
$\Omega^{1,1}_{N2-O}, \Omega^{2,2}_{N2-O}$	Epistemic	$\pm 20\%$
$\Omega^{1,1}_{N2-O2}, \Omega^{2,2}_{N2-O2}$	Epistemic	$\pm 20\%$

through sampling at a fraction of the cost, requiring just over 40 function/gradient evaluations for the two optimizations. Because the optimization results represent function values achieved using epistemic variables in the specified interval, the bounds produced from optimization should be viewed as the correct results. As previous results have shown, it is likely that further sampling would give results that approach the optimization bounds [35].

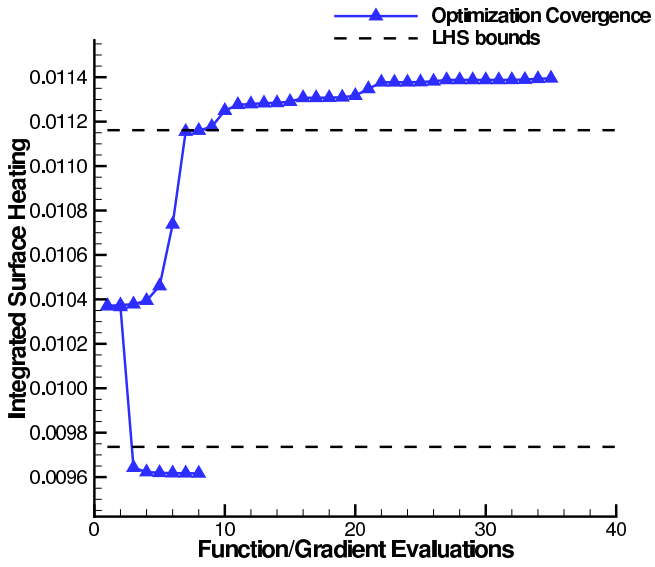

Figure 6. Convergence of Optimization over epistemic variables for fixed aleatory variables compared with bounds from sampling.

1.4.3. Statistics-of-Intervals/Kriging for Mixed Aleatory/Epistemic Uncertainty

To quantify uncertainty for the mixed case, a combined optimization/Kriging approach is applied in a statistics-of-intervals manner, meaning that multiple optimizations are performed over different values of the aleatory variables and a surrogate is created for these optimization results. To test this strategy, ten variables are treated as uncertainty, given in Table 7. The epistemic variables for the mixed case are the same eight variables as those used for the pure epistemic case. Two aleatory variables are used for this study, corresponding to the freestream density and velocity. These parameters and associated uncertainty were chosen based on the results of similar uncertainty quantification studies[5, 20]. The choice of which parameters to treat as aleatory was taken from an additional study[21].

The goal of a mixed aleatory/epistemic uncertainty quantification is to produce an interval containing the solution with a specified probability, known as a P-box [3]. The performance of the statistics-of-interval

approach is measured based on its predictions of a 99% P-box as well as its ability to predict the CDF curves associated with the minimum and maximum values of the optimization, enabling any P-box to be predicting in principle. These predictions can be validated in one of two ways. First, a nested sampling approach can be used, in which an aleatory uncertainty quantification is performed for each epistemic sample [3]. For the variables used in this study, this nested approach would require approximately 30 million simulation results. Aside from nested sampling, the statistics associated with the interval bounds can be validated by exhaustive sampling of the optimization results over the aleatory variables. Hence, thousands of pairs of optimization results would be required to validate the statistics associated with the interval bounds. For this test, approximately 300,000 simulation results would be required. For both approaches, this validation is beyond the computational budget of this work. To validate the SOI-kriging method applied to a real gas simulation, each element of the method was validated separately against exhaustive sampling. With each element validated, the mixed aleatory/epistemic uncertainty was calculated by using successively more accurate surrogate models to demonstrate convergence of the statistic predictions.

The optimization portion of the SOI method was validated in the previous section. With the optimization confirmed, the ability of the Kriging model to capture the aleatory variation of the integrated surface heating was tested. Although the previously discussed method for aleatory uncertainty quantification was based on a gradient-enhanced Kriging model, the SOI method uses a function-only Kriging model. Because the SOI method requires the construction of a Kriging model for the optimization results, the gradient of the optimal results would be required for a gradient-enhanced method, a quantity that is difficult to calculate. Luckily, within hypersonic flows, the number of aleatory variables is relatively small, reducing the need for a gradient-enhanced surrogate model.

In order to determine the number of training points required to capture the design space associated with the two variables identified in Table 7, Kriging models were constructed using an increasing number of simulation results and statistics were predicted based on each model. For this test, the epistemic variables were frozen at their non-perturbed values (1 in the terms of the parameters defined in Table 7), and sampling was performed over the aleatory variables. In order to provide validation data, Monte Carlo sampling was performed over the two aleatory variables, and the distribution was characterized both by constructing a CDF curve and by calculating specific statistics. In order to acquire accurate statistics, 4,564 samples were used, and a separate simulation was performed for each. With the validation data acquired, ordinary (constant mean function) kriging models with increasing numbers of training points were constructed. Because the epistemic variables for this test were fixed, each training point required only a single CFD simulation.

As a first test, the convergence of the mean, variance, and $99th$ percentile are shown for kriging models with increasing numbers of training points. The convergence of this metric as a function of training point number is given in Table 8. As the results show, predictions of the kriging model rapidly converge toward the Monte Carlo results. In addition to predicting distribution statistics, a CDF of the output is constructed based on samples extracted from the kriging model and compared with that of Monte Carlo sampling. Figure 7 shows the predicted CDF curve for a kriging model with 8 training points and the CDF from Monte Carlo sampling. Using only 8 samples, the kriging model produces a CDF curve nearly identical to the curve produced through Monte Carlo sampling, at a fraction of the cost. Based on this result, it can be expected that a similar number of optimization results should be required to accurately predict the mixed result. Since this problem only considers the uncertainty due to two aleatory variables, this cost is most likely overly optimistic for typical problems.

With each element of the SOI-kriging approach validated independently, the complete mixed aleatory/epistemic uncertainty is predicted by using optimization for the epistemic dependence and an ordinary kriging model for the aleatory dependence. In order to demonstrate the validity of the full results, the convergence of the minimum and maximum $99th$ percentile predictions are shown as the number of training points for the kriging model is increased. For the mixed results, a training point now represents a pair of optimizations and has a cost of approximately 60 function/gradient evaluations on average. Table 9 shows the convergence of the maximum $99th$ percentile and minimum $99th$ percentile as the number of training points is increased. As the table demonstrates, the statistic predictions quickly converge to asymptotic values. Included in Table 9 is the total cost in terms of function/gradient evaluations. While the nested sampling and exhaustive sampling of the optimization were prohibitively expensive for the CFD model, the SOI-kriging

Table 8. Convergence of Kriging Statistic Predictions for Aleatory Uncertainty with Fixed Epistemic Variables with Increasing Number of Training Points

Training Points	Average	Variance	99*th* Percentile
8	1.036110×10^{-2}	6.061055×10^{-8}	1.098518×10^{-2}
16	1.036622×10^{-2}	6.075630×10^{-8}	1.097558×10^{-2}
31	1.034997×10^{-2}	6.145065×10^{-8}	1.098506×10^{-2}
59	1.037171×10^{-2}	6.185576×10^{-8}	1.099184×10^{-2}
121	1.036669×10^{-2}	6.120957×10^{-8}	1.097695×10^{-2}

Monte Carlo Results

Samples	MC Average	MC Variance	MC 99*th* Percentile
4564	1.036082×10^{-2}	6.103365×10^{-8}	1.098385×10^{-2}

Table 9. 99*th* Percentile Predictions for SOI Method Using Ordinary Kriging Model for Real Gas CFD Simulation

Training Data Size	Number of F/G Evaluations	99^{th} percentile of Min	99^{th} percentile of Max
8	500	1.017556×10^{-2}	1.206949×10^{-2}
15	900	1.016681×10^{-2}	1.207132×10^{-2}
23	1400	1.018928×10^{-2}	1.207939×10^{-2}
52	3000	1.020232×10^{-2}	1.210513×10^{-2}
104	6176	1.020243×10^{-2}	1.210416×10^{-2}

model was able to capture converged statistics with a number of function/gradient evaluations within the computational budget (although still most likely prohibitively high for complex simulations). Nevertheless, by using the kriging model combined with optimization, the SOI-kriging method was able to quantify the mixed aleatory/epistemic uncertainty problem where other methods could not be used.

Figure 8 shows the convergence of the average and variance prediction based on kriging models with increasing numbers of training points. As this Figure shows, the statistics produced from each model show little variation as the number of training points increases and the variation is small compared with the overall interval produced due to the epistemic uncertainty. In addition to calculating specific statistics of the output interval, the CDF of the minimum and maximum values can be predicted by sampling from the kriging surface. The bounding CDF curves are plotted in Figure 9 for a kriging model based on 8 and 104 pairs of optimizations. As the figure demonstrates, the CDF curves are nearly identical, suggesting that the kriging model has reached some level of convergence.

1.5. Conclusions and Future Work

In this paper, gradient-based strategies for rapidly quantifying the uncertainty within hypersonic flows have been explored. In particular, the variability of an objective based on uncertain model parameters within the simulation were calculated. For aleatory uncertainties, a combined Kriging and regression model is used to represent the variation of the simulation output over the input variables, providing a basis for inexpensive Monte Carlo sampling. By incorporating derivative observations into this surrogate and a dimension reduction based on regression-based global sensitivity analysis, the model is applicable to problems containing a large number of variables, with a cost increasing linearly as the dimension increases. For epistemic uncertainties, gradient-based optimization is employed to determine the minimum and maximum output values possible given input parameters within specified ranges. In particular, the L-BFGS optimization algorithm is used for these minimization and maximization problems, leading to a method particularly suited for large dimensional problems. Finally, the problem of mixed aleatory/epistemic uncertainty is addressed via a mixed optimization/surrogate method where by multiple optimizations are performed for different values of the aleatory variables and a surrogate is constructed over these optimization results to determine the necessary

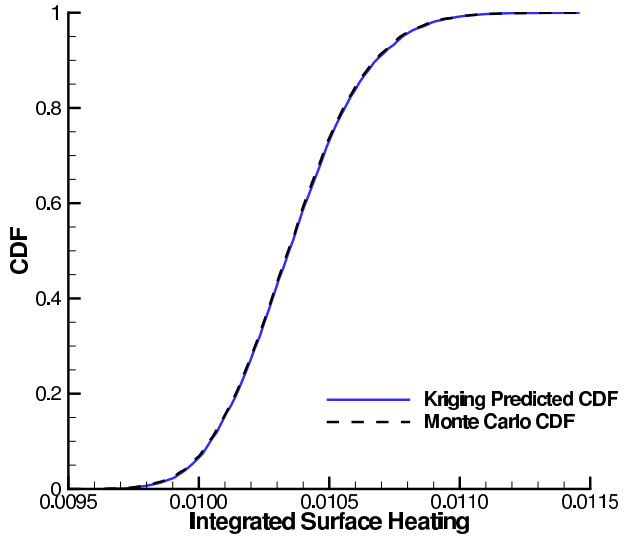

Figure 7. CDF based on Kriging model using 8 sample points compared with CDF of Monte Carlo results with fixed epistemic variables.

statistics of this interval. All of these methods were demonstrated based on a representative test case from a five species, two temperature non-equilibrium CFD simulation. For the case of aleatory and epistemic uncertainties, the proposed methods produced uncertainty predictions in good agreement with those from exhaustive sampling at a fraction of the cost. For the case of mixed uncertainties, acquiring the necessary validation data was prohibitively expensive, making the accelerated methods the only viable strategy for calculating this type of uncertainty.

For future work, the strategies presented here should be extended to a wider range of problems and parameters. To improve upon the surrogates used for the aleatory uncertainty, higher-order derivatives, such as the Hessian, can be incorporated into the training of the surrogate [22]. Additionally, more sophisticated techniques for determining the training data used within the construction of the surrogate can be utilized to increase accuracy without a corresponding increase in cost [51]. To improve upon the optimization-based methods, the optimization algorithm itself should be examined. Because the problem of epistemic uncertainty requires global extrema, efficient global optimization techniques should be explored within this context. Ideally, the efficient techniques should scale to large input dimension and leverage the additional information provided by gradient values. For aerospace applications, Kriging-based global optimization has been explored [22, 10, 11]. Although these surrogate-based methods may have difficulty extending to high dimension, these techniques combined with a suitable dimension reduction technique may provide a strong basis for epistemic uncertainty quantification in hypersonic flows. Finally, for situations in which gradient-based optimization appears sufficient for determining the global extrema, as was the case for the problem presented in this work, future work should focus on reducing the cost associated with this optimization, through the examination of more sophisticated optimization algorithms and the incorporation of Hessian information [52].

Acknowledgement

This work was supported by the U.S. Department of Energy through a Computational Science Graduate Fellowship under grant number DE-FG02-97ER25308 (Brian Lockwood). The contributions of other collaborators including Mihai Anitescu, Markus Rumpfkeil and Wataru Yamazaki are also acknowledged and

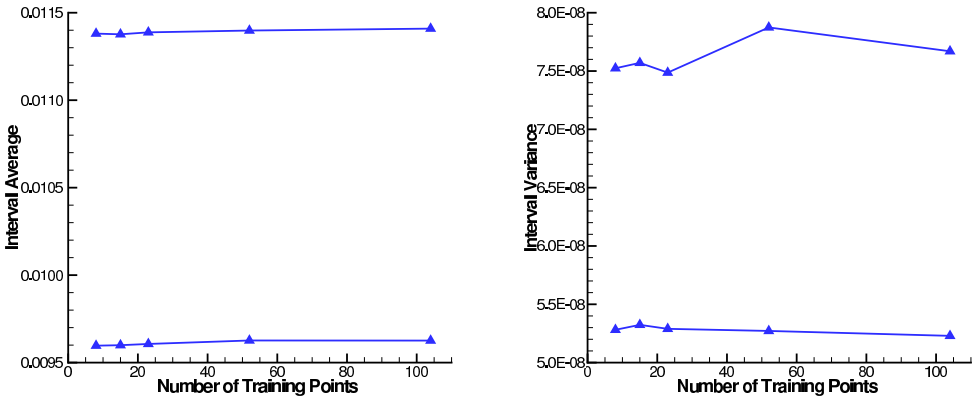

Figure 8. Convergence of average (Left) and variance (Right) prediction for minimum and maximum distribution using kriging models built from increasing numbers of optimization results for real gas CFD simulation.

greatly appreciated.

References

[1] J. M. Luckring, M. J. Hemsch, J. H. Morrison, Uncertainty in computational aerodynamics, in: 41st AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 2003, AIAA Paper, 2003-0409.
[2] C. R. Gumbert, P. A. Newman, G. J. Hou, Effect of random geometric uncertainty on the computational design of 3-d wing, in: 20th AIAA Applied Aerodynamics Conference, St. Louis, MO, 2002, AIAA Paper, 2002-2806.
[3] C. J. Roy, W. L. Oberkampf, A complete framework for verification, validation and uncertainty quantification in scientific computing, in: 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2010, AIAA Paper, 2010-124.
[4] M. J. Wright, D. Bose, Y.-K. Chen, Probabilistic modeling of aerothermal and thermal protection material response uncertainties, AIAA Journal 45 (2) (2007) 399–425.
[5] G. E. Palmer, Uncertainty analysis of cev leo and lunar return entries, in: 39th AIAA Thermophysics Conference, Miami, FL, 2007, AIAA Paper, 2007-4253.
[6] D. Ghate, M. Giles, Inexpensive Monte Carlo uncertainty analysis, in: B. Uthup, S. Koruthu, R. Sharma, P. Priyadarshi (Eds.), Recent Trends in Aerospace Design and Optimization, Tata McGraw-Hill, New Delhi, 2006, pp. 203–210.
[7] D. P. Ghate, M. B. Giles, Efficient hessian calculation using automatic differentiation, in: 25th AIAA Applied Aerodynamics Conference, Miami, FL, 2007, AIAA Paper, 2007-4059.
[8] N. Cressie, The Origins of Kriging, Mathematical Geology 22 (3) (1990) 239–252.
[9] J. R. Koehler, A. B. Owen, Computer Experiments, in: Handbook of Statistics, pp. 261-308, 1996.
[10] D. R. Jones, M. Schonlau, W. J. Welch, Efficient Global Optimization of Expensive Black-Box Functions, Journal of Global Optimization 13 (1998) 455–492.
[11] T. W. Simpson, J. J. Korte, T. M. Mauery, F. Mistree, Comparison of response surface and kriging models for multidisciplinary design optimization, in: 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, 1998, AIAA Paper, 98-4758.
[12] H. S. Chung, J. J. Alonso, Using gradients to construct cokriging approximation models for high-dimensional design optimization problems, in: 40th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 2002, AIAA Paper, 2002-0317.
[13] J. D. Martin, T. W. Simpson, Use of Kriging Models to Approximate Deterministic Computer Models, AIAA Journal 43 (4) (2005) 853–863,.
[14] S. Jeong, M. Murayama, K. Yamamoto, Efficient Optimization Design Method Using Kriging Model, Journal of Aircraft 42 (No. 2) (2005) 413–420.
[15] J. Peter, M. Marcelet, Comparison of Surrogate Models for Turbomachinery Design, WSEAS Transactions on Fluid Mechanics 3 (1) (2008) 10–17.
[16] J. Laurenceau, P. Sagaut, Building Efficient Response Surfaces of Aerodynamic Functions with Kriging and Cokriging, AIAA Journal 46 (2) (2008) 498–507.
[17] J. Laurenceau, M. Meaux, Comparison of gradient and response surface based optimization frameworks using adjoint method, in: 49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Schaumburg, IL, 2008, AIAA Paper, 2008-1889.
[18] W. Yamazaki, S. Mouton, G. Carrier, Efficient design optimization by physics-based direct manipulation free-form deformation, in: 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Victoria, Canada, 2008, AIAA Paper, 2008-5953.

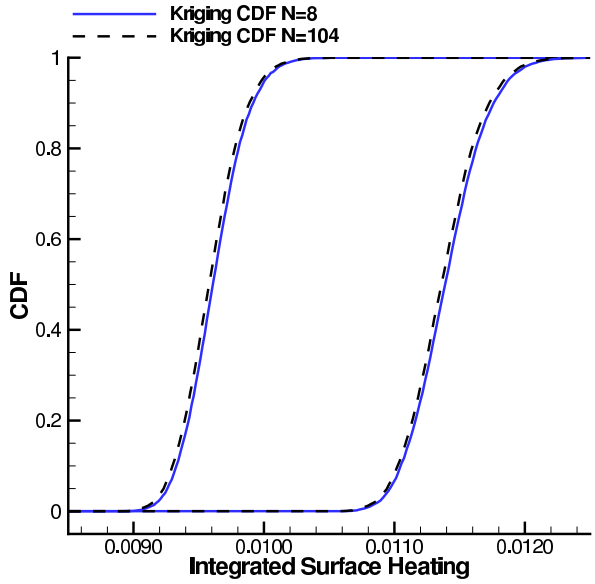

Figure 9. Kriging-predicted CDF curves for maximum and minimum values using 8 and 104 optimization pairs.

[19] B. A. Lockwood, M. P. Rumpfkeil, W. Yamazaki, D. J. Mavriplis, Uncertainty quantification in viscous hypersonic flows using gradient information and surrogate modeling, in: 49th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2011, AIAA Paper, 2011-885.

[20] A. Alexeenko, A. Weaver, R. Greendyke, J. Camberos, Flowfield uncertainty analysis for hypersonic cfd simulations, in: 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2010, AIAA Paper, 2010-1180.

[21] B. R. Bettis, S. Hosder, Uncertainty quantification in hypersonic reentry flows due to aleatory and epistemic uncertainties, in: 49th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2011, AIAA Paper, 2011-252.

[22] W. Yamazaki, M. P. Rumpfkeil, D. J. Mavriplis, Design optimization utilizing gradient/hessian enhanced surrogate model, in: 40th Fluid Dynamics Conference and Exhibit, Chicago, IL, 2010, AIAA Paper, 2010-4363.

[23] O. Roderick, M. Anitescu, P. Fischer, Polynomial Regression Approaches Using Derivative Information for Uncertainty Quantification, Nuclear Science and Engineering 164 (2) (2010) 122–139.

[24] B. A. Lockwood, M. Anitescu, Gradient-enhanced universal kriging for uncertainty propagation in nuclear engineering, Nuclear Science and Engineering 170 (2) (2012) 168–195.

[25] O. Pironneau, On Optimum Design in Fluid Mechanics, Journal of Fluid Mechanics 64, No. 1 (1974) 97–110.

[26] A. Jameson, Optimum Aerodynamic Design Using Control Theory, in: O. K. Hafez, M. (Ed.), Computational Fluid Dynamics Review, Wiley: New York, 1995, pp. 495–528.

[27] R. M. Errico, What is an adjoint model ?, Bulletin of the American Meteorological Society 8 (11) (1997) 2577–2591.

[28] B. A. Lockwood, D. J. Mavriplis, Parameter sensitivity analysis for hypersonic viscous flow using a discrete adjoint approach, in: 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2010, AIAA Paper, 2010-447.

[29] W. L. Kleb, C. O. Johnston, Uncertainty analysis of air radiation for lunar return shock layers, in: AIAA Atmospheric Flight Mechanics Conference and Exhibit, Honolulu, HI, 2008, AIAA Paper, 2008-6388.

[30] J. C. Helton, J. D. Johnson, W. L. Oberkampf, C. J. Sallaberry, Representation of Analysis Results Involving Aleatory and Epistemic Uncertainty, Tech. Rep. SAND 2008-4379, Sandia National Laboratories (2008).

[31] R. R. Yager, L. Liu, Classic Works of the Dempster-Shafer Theory of Belief Functions. Studies in Fuzziness and Soft Computing Series, Vol. 219, Berlin: Springer, 2008.

[32] V. Kreinovich, S. Ferson, A new cauchy-based black-box technique for uncertainty in risk analysis, in: Reliability Engineering and Systems Safety, 2002, pp. 267–279.

[33] M. Pilch, T. G. Trucano, J. C. Helton, Ideas Underlying Quantification of Margins and Uncertainties (QMU): A white paper, Tech. Rep. SAND2006-5001, Sandia National Laboratories (2006).

[34] M. S. Eldred, L. P. Swiler, Efficient algorithms for mixed aleatory-epistemic uncertainty quantification with application to radiation-hardened electronics, Tech. Rep. SAND2009-5805, Sandia National Laboratories (2009).

[35] B. A. Lockwood, M. Anitescu, D. J. Mavriplis, Mixed aleatory/epistemic uncertainty quantification for hypersonic flows via gradient-based optimization and surrogate models, in: 50th AIAA Aerospace Sciences Meeting and Exhibit, Nashville, TN, 2012, AIAA Paper, 2012-1254.

[36] B. Hassan, Thermo-chemical nonequilibrium effects on the aerothermodynamics of hypersonic vehicles, Ph.D. thesis, North Carolina State University, Albuquerque, NM (December 1993).

[37] P. A. Gnoffo, R. N. Gupta, J. L. Shinn, Conservation equations and physical models for hypersonic air flows in thermal and chemical nonequilibrium, Tech. rep., NASA (February 1989).
[38] D. R. Olynick, A new lu-sgs flow solver for calculating reentry flows, Ph.D. thesis, North Carolina State University (1992).
[39] J. R. Edwards, A low-diffusion flux-splitting scheme for navier-stokes calculations, Computers & Fluids 26 (6) (1997) 635–659.
[40] M.-S. Liou, A sequel to ausm, part ii: Ausm+-up for all speeds, Journal of Computational Physics 214 (2006) 137–170.
[41] F. M. Cheatwood, P. A. Gnoffo, User's manual for the langley aerothermodynamic upwind relaxation algorithm (laura), Tech. rep., NASA (April 1996).
[42] NASA, FUN3D: Fully Unstructured Navier-Stokes Manual, http://fun3d.larc.nasa.gov/index.html (May 2009).
[43] D. Mavriplis, Multigrid solution of the discrete adjoint for optimization problems on unstructured meshes, AIAA Journal 44 (1) (2006) 42–50.
[44] L. Hascoët, TAPENADE: A Tool for Automatic Differentiation of Programs, in: Proceedings of 4^{th} European Congress on Computational Methods, ECCOMAS'2004, Jyvaskyla, Finland, 2004.
[45] S. Hosder, R. W. Walters, Non-intrusive polynomial chaos methods for uncertainty quantification in fluid dynamics, in: 48th AIAA Aerospace Sciences Meeting, Orlando, IL, 2010, AIAA Paper, 2010-129.
[46] Z.-H. Han, S. Görtz, R. Zimmermann, On improving efficiency and accuracy of variable-fidelity surrogate modeling in aero-data for loads context, in: CEAS 2009 European Air and Space Conference, Manchester, UK, 2009.
[47] C. Rasmussen, C. Williams, Gaussian Processes for Machine Learning, The MIT Press, 2006.
[48] W. L. C. R. Ercan Solak, Roderick Murray-Smith, D. Leith, Derivative observations in gaussian process models of dynamic systems (2003).
[49] M. P. Rumpfkeil, W. Yamazaki, D. J. Mavriplis, Uncertainty analysis utilizing gradient and hessian information, in: Sixth International Conference on Computational Fluid Dynamics, ICCFD6, St. Petersburg, Russia, 2010.
[50] C. Zhu, R. H. Byrd, P. Lu, J. Nocedal, L-BFGS-B: A Limited Memory FORTRAN Code for Solving Bound Constrained Optimization Problems, Tech. Rep. NAM-11, Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois, USA (1994).
[51] M. P. Rumpfkeil, W. Yamazaki, D. J. Mavriplis, A dynamic sampling method for kriging and cokriging surrogate models, in: 49th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2011, AIAA Paper, 2011-883.
[52] M. P. Rumpfkeil, D. J. Mavriplis, Efficient Hessian Calculations using Automatic Differentiation and the Adjoint Method with Applications, AIAA Journal 48 (10).