Mixed Aleatory/Epistemic Uncertainty Quantification for Hypersonic Flows via Gradient-Based Optimization and Surrogate Models*

Brian A. Lockwood[†]

University of Wyoming, Laramie, WY, 82071, USA

Mihai Anitescu[‡]

Argonne National Laboratory, Argonne, IL, 60439, USA

and

Dimitri J. Mavriplis[§]

University of Wyoming, Laramie, WY, 82071, USA

The use of optimization for the propagation of mixed epistemic/aleatory uncertainties is demonstrated within the context of hypersonic flows. Specifically, this work focuses on strategies applicable for models where input parameters can be divided into a set of variables containing only aleatory uncertainties and a set with epistemic uncertainties. With the input parameters divided in this way, uncertainty due to the epistemic variables is propagated via a constrained optimization approach, while the uncertainty due to aleatory variables is propagated via sampling. A statistics-of-intervals approach is proposed in which the constrained optimization results are treated as a random variable and multiple optimizations are performed to quantify the aleatory uncertainty. In order to reduce the total number of optimizations required, a surrogate is employed to model the variation of the optimization results with respect to the aleatory variables, and exhaustive sampling is performed on this surrogate to determine the desired statistics. The properties of the statistics-of-intervals approach are demonstrated by using the Fay-Riddell stagnation heating correlations and compared with a competing method based on uncertain optimization. Additionally, the statistics-of-intervals approach is demonstrated for mixed epistemic/aleatory uncertainty quantification on a real gas computational fluid dynamic simulation.

I. Introduction and Motivation

As the role of computational modeling in the design and analysis of complex engineering systems has grown, uncertainty quantification has become an increasingly important part of computational science, providing the ability to assess the quality of computational results and apply confidence bounds to output metrics. Uncertainty quantification is of particular importance for problems in which experimental data is difficult or impossible to obtain, as is the case with hypersonic flows. The uncertainty in simulation outputs can arise from various sources, 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. Although a number of techniques exist for the propagation of aleatory uncertainties, few are readily extended to epistemic or mixed uncertainties. Nevertheless, regulatory agencies and design teams are increasingly being asked to specifically characterize and

*Preprint ANL/MCS-P1988-1211, Argonne National Laboratory, Mathematics and Computer Science Division

[†]Graduate Student, Department of Mechanical Engineering; blockwoo@uwyo.edu, Student Member AIAA

[‡]Computational Mathematician, Mathematics and Computer Science Division, Argonne National Laboratory

[§]Professor, Department of Mechanical Engineering; mavripl@uwyo.edu, Associate Fellow AIAA

quantify epistemic uncertainty in conjunction with aleatory uncertainties.³ In this paper, a method for the quantification of uncertainties arising from both epistemic and aleatory sources will be presented within the context of hypersonic flows.

Aleatory uncertainties arise due to the inherent randomness of a variable and are characterized by a probability distribution.⁴ 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. Monte Carlo sampling has been used within the context of hypersonic flows to build up the required statistics for relevant simulation outputs.^{5,6} 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^{7,8} 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^{9–19} although their use within the context of uncertainty quantification for hypersonic flows is limited.²⁰ Techniques based on polynomial chaos have also been employed with success in the context of hypersonic flows.^{21,22} 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.^{13,17,18,23–25} When adjoint methods are employed, this gradient may be evaluated with a cost approximately equal to the simulation of the physical problem.^{26–28} By incorporating derivative values, the cost associated with training an accurate surrogate can be greatly reduced.

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 uncertainties present in hypersonic flows and previous studies assuming pure aleatory uncertainties, although important initial steps, have likely underestimated the uncertainty associated with simulation objectives.^{5,29} 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.^{30–32} 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.⁴ 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 (or second-order probability). 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.^{4,34} 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 using sampling of 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.⁴ 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 and, when combined with gradient-enhancement, could be applied to problems of moderate dimension.²² 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.⁴ Stated in other terms, the bounds of the confidence interval of the output distribution must itself 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 (referred to as uncertainty optimization, or UQOPT, throughout) 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.³⁴ For each step in the optimization, the aleatory uncertainty is quantified, and the relevant statistic of the distribution is calculated and used as the metric for the optimization.

In this paper, a statistics-of-the-interval approach (SOI for short) is proposed within the context of hypersonic flows. For this approach, an optimization problem will be posed for each set of aleatory variables, and repeated optimizations will be used to determine the relevant statistics of the interval. Because of the adjoint capability of the flow solver used for this work, a gradient-based optimization is used, reducing the cost of each optimization and ensuring optimal scaling as the input dimension increases. 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.

The structure of this paper is as follows. In Section II, details of the physical problems and flow solver used for this work are presented. Because the availability of gradient information is vital to this work, Section III outlines the discrete adjoint implementation and sample sensitivity results are shown. In Section IV, details of the proposed strategy are given and compared with the competing uncertain optimization approach. Demonstration results using the models of section II are given in Section V.

II. Flow Problem

This section details the physical models used to test the previously described uncertainty quantification strategies. Two models were used for this work: the Fay-Riddell stagnation heating correlation and a non-equilibrium real gas CFD solver.

II.A. Fay-Riddell Heating

In order to provide a simplified, computationally inexpensive test problem that can mimic the characteristics of hypersonic CFD simulations with respect to model parameters and design space complexity, the Fay-Riddell stagnation heating correlation is examined.³⁵ This model is an analytic function, permitting the use of exhaustive sampling to validate the uncertainty quantification strategies examined in this paper. The equations for this model are as follows.

$$q'' = 0.76(Pr_w)^{-0.6}(\rho_w\mu_w)^{0.1}(\rho_e\mu_e)^{0.4}\sqrt{\left(\frac{dU_e}{dx}\right)}(h_{o,e} - h_w)\left[1 + (Le^{0.52} - 1)\left(\frac{h_D}{h_{o,e}}\right)\right] \quad (1)$$

$$\left(\frac{dU_e}{dx}\right) = \frac{1}{R_N}\sqrt{2\frac{p_e - p_\infty}{\rho_e}} \quad (2)$$

$$h_D = \sum_i C_{i,e}\Delta h_{f,i}^o \quad (3)$$

In these equations, q'' is the heat flux at the stagnation point, Pr represents the Prandtl number, Le is the Lewis number, R_N is the radius of curvature, $\Delta h_{f,i}^o$ is the heat of formation for each species, h is the internal enthalpy and h_o is the total enthalpy. For this equation, the subscript e represents properties at the edge of

the boundary layer, and w represents values at the wall. The properties at the edge of the boundary layer are found by solving the normal shock problem using equilibrium properties for air. From the curve fits for the equilibrium properties of air, the enthalpy downstream of the shock is first written in terms of the density and pressure. A root-finder is then used to determine the downstream density and pressure required to satisfy the Rankine-Hugoniot jump conditions. With this density and pressure determined, all other bulk properties are calculated via curve fits.³⁶ The composition after the shock is determined via statistical thermodynamics. Through analytic expressions of each species partition function, the equilibrium constant for each dissociation reaction is calculated. These equilibrium constants are then related to the partial pressure of each species and a Newton solver is used to solve for the partial pressure of each species.^{37,38} Using these partial pressures, the mass fraction for each species can be found. Further details of the solution process used for the Fay-Riddell heating correlation can be found in other work.³⁹ The required viscosities and thermal conductivity are evaluated by using a collision integral transport model that is common to the CFD solver and the Fay-Riddell problem.^{40,41} This model was chosen in order to maintain the same input parameters between the Fay-Riddell model and the real gas CFD code. The Prandtl number at the wall is computed based on the transport values evaluated at the wall, and the Lewis number is evaluated based on the Prandtl number and assuming a constant Schmidt number of 0.5.⁴²

II.B. Real Gas 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

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \vec{F}(\mathbf{U}) = \nabla \cdot \vec{F}_v(\mathbf{U}) + \mathbf{S}(\mathbf{U}), \quad (4)$$

where $\mathbf{U}$ are the conserved flow variables, $\vec{F}$ is the inviscid flux, $\vec{F}_v$ is the viscous flux, and $\mathbf{S}$ 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.⁴² Both the Dunn-Kang chemical kinetics model and the Park model have been implemented for this work. Despite the superior accuracy of the Park model, the Dunn-Kang model is generally used throughout this work because of the ease with which uncertain parameters within the model may be specified. 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.⁴⁰ 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.^{40,41} 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:

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

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.⁴¹⁻⁴³ 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.^{43,44}

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. Middle: Surface heating distribution. Right: Comparison of temperatures along centerline with LAURA⁴⁵ 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^{45,46} and are depicted in Figure 1. 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.

III. Adjoint Sensitivity Analysis

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.⁴⁷ 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.

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

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.

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

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.⁴⁸

$$\frac{dL}{dD} = \frac{\partial L}{\partial D} + \frac{\partial L}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial D} \quad (8)$$

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 $\mathbf{U}$.

$$\frac{\partial \mathbf{R}}{\partial D} + \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial D} = 0 \quad (9)$$

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).

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

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

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

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

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

A sample adjoint solution for the 5 km/s benchmark is found in Figure 2. 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.

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

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.⁴⁸ 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.⁴⁹ 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.

In order to demonstrate the use of the adjoint, the sensitivities of surface heating with respect to parameters governing the chemical kinetics model and transport coefficients were calculated and are presented for the 5 km/s real gas benchmark. The objective used for these demonstration results is integrated surface heating, given by the following equation.

$$L = -\frac{\int_{\partial\Omega} k \nabla T \cdot \vec{n} + k_v \nabla T_v \cdot \vec{n} dA}{\frac{1}{2} \rho_\infty V_\infty^3} \quad (14)$$

In this equation, T is the translational-rotational temperature, k is the translational-rotational thermal conductivity, T_v is the vibrational temperature, and k_v is the vibrational thermal conductivity. The first

Figure 2. Density adjoint for integrated surface heating

variables examined relate to the specification of reaction rates. For the Dunn-Kang chemical kinetics model used within this work, the reaction rates take the following form.

$$K_f = C_f T_a^{\eta_f} e^{-\frac{E_{a,f}}{k_B T_a}} \quad (15)$$

$$K_b = C_b T_a^{\eta_b} e^{-\frac{E_{a,b}}{k_B T_a}} \quad (16)$$

Here $E_{a,f}$ 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 parameters examined in this case were C_f and C_b for each reaction, giving a total of 34 parameters. Figure 3 depicts the computed sensitivity of surface heating with respect to the forward and backward reaction rates using the adjoint procedure for the 5 km/s benchmark case. Note that because of the large discrepancy between the design variable and the objective, the sensitivity is expressed as fractional change in objective per fractional change in design variable (i.e. $\frac{dL}{dD}$).⁴⁷ As the results demonstrate, the reactions governing the production and breakdown of NO , as well as the oxygen recombination reactions, have the greatest influence on integrated surface heating.

In addition to reaction rates, the sensitivity with respect to parameters within the transport model was calculated. For the collision integral model, measured collision integrals between the five species at 2000 K and 4000 K are used, and linear interpolation is used to determine the collision integral at the appropriate temperature.⁴⁰

$$\log_{10}(\Omega_{s,r}^{k,k}) = \log_{10}(\Omega_{s,r}^{k,k})_{2000} + [\log_{10}(\Omega_{s,r}^{k,k})_{4000} - \log_{10}(\Omega_{s,r}^{k,k})_{2000}] \frac{\ln(T) - \ln(2000)}{\ln(4000) - \ln(2000)} \quad (17)$$

In the above equation, $\Omega_{s,r}^{k,k}$ represents the collision integral between species s and species r . For the five-species model, 15 independent collision interactions are possible. This model gives 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. In reality, it is likely that the uncertainty for the collision integrals measured at 2000 K and 4000 K are correlated. Because of this assumption, the sensitivities with respect to only the collision integrals at 2000 K were calculated. The results of this calculation are also presented in Figure 3. As the results show, the collisions involving N_2 have the greatest effect on integrated surface heating. This result is unsurprising because N_2 is the predominant species at the super-catalytic wall boundary condition.

where $m(D)$ is the mean function of the distribution and K is the covariance matrix for the data. In the context of the SOI method, the variable y represents either the minimum (L_{min}) or maximum (L_{max}) value from the optimization results. For this work, the mean function is typically assumed constant $m(D) = \mu_0$, corresponding to ordinary kriging. Higher-order polynomials, such as Hermite or Legendre polynomials, may also be used in the kriging model to enhance accuracy. When a polynomial basis is used in the kriging model, the approach is referred to as universal 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 $5/2$ since this combination of function and smoothness parameter produced the most accurate results.^{50, 51} 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 of y away from these observations is predicted for an ordinary kriging model using the following equation.

$$y^* = \mu_0 + k(D^*, D)K^{-1}(\mathbf{y} - \mu_0) \quad (22)$$

Here, $k(D^*, D)$ is the covariance between the test point of interest D^* and the training data D . Although gradient or Hessian information can be easily incorporated into the training of a Kriging model, the model explored here is based on function values only because of the relatively small number of aleatory variables.^{23, 50}

Because the SOI approach does not require differentiation of the surrogate model, sophisticated surrogate prediction techniques can be used in conjunction with the method. Because of the stochastic nature of kriging models, in addition to predicting mean behavior, the distribution of model predictions can be constructed by repeated sampling of the Gaussian process representing the kriging model. This distribution is due to the uncertainty inherent in the surrogate model and can provide an error estimate for the surrogate predictions. In order to avoid confusion, results based on sampling from the kriging model's Gaussian process are not presented; however, this approach is presented for the case of purely aleatory uncertainty in reference.²⁵

IV.B. Uncertain Optimization Approach

In addition to the proposed method, an uncertain optimization (UQOPT) approach is considered.³⁴ In this approach, a single pair of optimization problems (one minimization and maximization) is performed over the epistemic variables. The objective for this optimization is a statistical quantity due to the aleatory uncertainty present in the problem. This process is represented by using the following equations. Within these equations, β again represents the set of epistemic variables, and α represents the aleatory variables. In contrast to the statistics-of-intervals approach, the statistical metric of interest must be chosen before the optimization is performed. Let the variable $J(\beta)$ represent a statistical metric based on the simulation output $L(\alpha, \beta)$, such as a variance or reliability metric. The optimization problem is now posed as follows.

$$J_{max} = \max_{\beta} J(\beta) \quad (23)$$

$$J_{min} = \min_{\beta} J(\beta) \quad (24)$$

An example of $J(\beta)$ is given by the reliability metrics seen in equations (32) and (35). In order to determine the metric $J(\beta)$, the distribution of $L(\alpha, \beta)$ based on variations in α must be characterized for each β encountered during the optimization process. In order to accelerate this process, a surrogate of $L(\alpha, \beta)$ with respect to α is created, and statistics are calculated based on this surrogate. Because of the ease with which statistics can be calculated in closed form, a polynomial chaos expansion is used to calculate $J(\beta)$, yielding the method denoted as UQOPT-PC.³⁴ This expansion is represented as follows.

$$L(\alpha, \beta) = \sum_{i=0}^{N_p} \hat{L}_i(\beta) \phi_i(\alpha) \quad (25)$$

Here, $\hat{L}$ are the expansion coefficients, and $\phi(\alpha)$ is a polynomial in the random variable. The exact choice of polynomial depends on the input distribution with Hermite polynomials used for normal distributions and Legendre polynomials for uniform distributions.⁵² Since the expansion polynomials are orthogonal, the determination of the expansion coefficients requires a simple inner product that can be computed by quadrature.

$$\hat{L}_i(\beta) = \int_R L(\alpha', \beta) \phi_i(\alpha') d\alpha = \sum_{k=0}^{N_q} w_k L(\alpha_k, \beta) \phi_i(\alpha_k) \quad (26)$$

Here, R is the support of the basis function ϕ , w_k are the quadrature weights, and N_q is the total number of quadrature points. Once the coefficients are determined, the mean and variance are given by the following.

$$L_\mu = \hat{L}_0 \quad (27)$$

$$L_{\sigma^2} = \sum_{i=1}^{N_p} \hat{L}_i^2 \int_R \phi_i(\alpha')^2 d\alpha' \quad (28)$$

Because of these explicit relations, the derivatives required for a gradient-based optimizer are easily computable for many statistical metrics. Because the optimization is performed over the epistemic variables, only the derivative with respect to β is required. This gradient is easily computed using the following equations.³⁴

$$\frac{\partial L(\alpha, \beta)}{\partial \beta} = \sum_{i=0}^{N_p} \frac{\partial \hat{L}_i(\beta)}{\partial \beta} \phi_i(\alpha) \quad (29)$$

$$\frac{\partial \hat{L}_i(\beta)}{\partial \beta} = \int_R \frac{\partial L(\alpha', \beta)}{\partial \beta} \phi_i(\alpha') d\alpha' = \sum_{k=0}^{N_q} w_k \frac{\partial L(\alpha_k, \beta)}{\partial \beta} \phi_i(\alpha_k) \quad (30)$$

Although technically any surrogate or statistical metric $J(\beta)$ can be used with this method, the desire to use a gradient-based optimizer limits the type of surrogate and statistical metric that can be used. In order to use a gradient-based optimizer, both the statistical metric and the surrogate process used to produce the metric must be differentiable, limiting the choice of statistical metric to distribution moments and choice of surrogate to deterministic models.

IV.C. Method Comparison

To compare the proposed statistics-of-intervals approach to the expected behavior of the uncertain optimization method, an analogy is first drawn between the latter method and another well-known method of optimization under uncertainty: stochastic programming.⁵³ A fairly encompassing formulation for the latter can be written as follows.

$$\min_{\beta} J(\beta) \equiv F(\beta, \mathbb{E}_{\alpha} G(L(\beta, \alpha))) \quad (31)$$

(Note that maximization and minimization problems are formally the same if the sign of the objective function is changed, so only minimization is discussed). Here G is a vector-valued function. Indeed, several of the well-known formulations can be such represented. For example, in decision theory set-ups, a risk function can be defined as $J(\beta) = \mathbb{E}_{\alpha} L(\beta, \alpha)$ (in this context, the function L is called the "loss function").⁵⁴ Another fairly common circumstance is that where the functional of interest depends on moments of a function of the random variable. For example (as will be used in Section §V.A.3) the functional of interest can be a combination of the mean and standard deviation of the random variable $L(\beta, \alpha)$, parametric in β .

$$J(\beta) = \mathbb{E}_{\alpha} L(\beta, \alpha) + c \sqrt{\mathbb{E}_{\alpha} L(\beta, \alpha)^2 - (\mathbb{E}_{\alpha} L(\beta, \alpha))^2} \quad (32)$$

Here, $\mathbb{E}_{\alpha}$ represents the expected value over the aleatory variables. If c is a given quantile of the standard normal, then $J(\beta)$ can be seen as an approximation of the quantile of the same rank of $J(\beta, \cdot)$ under a normal approximation. It can be immediately seen that either case can be abstractly set up as a functional defined by (31).

Therefore, many analytical conclusions concerning the comparison between the uncertain optimization approach and the statistics-of-intervals approach can be reached by reasoning on (31). An immediate contention is whether the risk function can be easily formulated. While this question may seem slightly irrelevant

to an algorithmically inclined scientist, it is a huge topic in decision theory, where there exists a massive literature on the elicitation of the risk function, for example.⁵⁴ It thus seems realistic that, particularly in the initial stages of the analysis, also known as the prospective stage, the proper objective function J may not be exactly known, except perhaps for a range of parameters (for example, in the case of (32), the parameter c may be known only within a factor of 2). In this case, the SOI approach is likely to be more efficient, as it will give a result which is easy to compute for a large class of objective functions, whereas the uncertain optimization would, to some extent, have to solve a new problem for each of the different objective functionals.

If the objective functional is known, the uncertain optimization method would seem to have an advantage in that it needs to run only one pair of optimization problems. Moreover, it is clear that F is not necessarily monotonic with respect to L in (31); hence the statistics-of-intervals approach may not actually solve (31). On the other hand, virtually all functionals of interest have this monotonicity property, so this situation is not as restrictive practically as it may appear. Moreover, the issue of “only one pair of optimization problems” is a bit superficial in that (31) will likely be approximated by sample average (SAA) approximation. That is, one constructs the i.i.d- α samples α_i , $i = 1, 2, \dots, n$, and constructs the SAA approximation:

$$\min_{\beta} J(\beta) \approx \min_{\beta} \hat{J}^N(\beta) \equiv F(\beta, \frac{1}{N} \sum_{i=1}^N G(L(\beta, \alpha_i))) \quad (33)$$

Therefore, in evaluating $\hat{J}^N$ a large number of α 's may need to be explored for each β , therefore putting the two methods on comparable footing, even for a well-defined J .

In addition to theoretical concerns, several practical differences exist between the two methods. The use of a gradient-based optimizer requires that the metric $J(\beta)$ used with the UQOPT method be differentiable. Additionally, the process used to inexpensively predict $J(\beta)$ based on a surrogate must also be differentiable. The statistics-of-intervals approach has no such limitations on the statistics that can be computed or the surrogate that can be used to model the aleatory variability. Because a number of independent optimizations are performed and surrogates are constructed to represent the variation of optimization results due to aleatory variables, the entire distribution of the interval is available provided the surrogate is of sufficient accuracy. Hence, the full empirical distribution function of each interval bound can be constructed, and any statistic may be calculated. Additionally, because the surrogate is applied merely to a collection of optimization sampling results, any surrogate may be used to represent the aleatory distribution. This flexibility may be important if the use of more sophisticated, stochastic surrogate models is desired.

V. Demonstration Results

In this section, the proposed statistics-of-intervals approach is demonstrated on problems relating to the hypersonic fluid flow models described in Section II. In order to prove the validity of the proposed method for mixed uncertainty problems and demonstrate desirable properties of the method, the Fay-Riddell heating correlation is first used. Because this correlation is an explicit function, exhaustive sampling in every aspect of the problem can be performed. In addition to showing the desirable properties of the statistics-of-intervals approach, the uncertain optimization method is demonstrated on the Fay-Riddell heating correlation and compared with the SOI approach. With the properties of the statistics-of-intervals approach demonstrated using the Fay-Riddell heating correlation, mixed aleatory/epistemic uncertainty is quantified for a real gas CFD simulation.

Throughout, the 5 km/s test case detailed in Section II is used. Ten uncertain parameters were chosen for all these tests and because of the set up of the Fay-Riddell heating correlation, both the CFD results and Fay-Riddell results use the same uncertain parameters. For all these tests, two input parameters are treated as aleatory and eight as epistemic. The aleatory variables corresponded to the freestream density and velocity, and the epistemic variables correspond to the eight most sensitive collision integrals from the sensitivity analysis presented in section III for the CFD results. The uncertain parameters along with the associated uncertainties are found in Table 2. These parameters and associated uncertainty were chosen based on the results of similar uncertainty quantification studies.^{6,21} The choice of which parameters to treat as aleatory was taken from an additional study.²²

The form of the uncertainty on the collision integrals is taken from an earlier study⁶ and assumes a multiplicative constant on Ω over the entire temperature range. This relationship can be represented as

Table 2. Uncertain Model Parameters

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

follows.

$$\Omega_{s,r}^{k,k} = A_{s,r} \hat{\Omega}_{s,r}^{k,k} \quad (34)$$

Here, $\hat{\Omega}_{s,r}^{k,k}$ is the mean collision integral between species s and r and $A_{s,r}$ is a variable that varies between 0.8 and 1.2 for the uncertainty specified in Table 2. Because the collision integrals use the base-10 logarithm of the collision integral interpolated between 2000 K and 4000 K, the cross-sections can be perturbed according to equation (34) by simply adding $\log_{10} A_{s,r}$ to the log of the mean collision integral terms evaluated at 2000 K and 4000 K.

For these tests, the uncertainty in surface heating is desired. For the Fay-Riddell heating correlation, the uncertainty in stagnation point surface heating due to the input uncertainties in Table 2 is predicted. For the CFD results, the uncertainty in integrated surface heating is desired, calculated based on Equation (14). The specification of the uncertainty is presented either through plots of the empirical distribution function of surface heating or based on a relevant statistical metric, such as the mean value reliability metric or a specified quantile.

V.A. Fay-Riddell Heating Results

In order to provide a baseline for the other methods illustrated in this work, exhaustive nested sampling was performed, and a series of cumulative distribution function (CDF) curves was constructed (known as a horse-tail plot). This plot is shown in Figure 4. For this baseline result, 5,000 aleatory samples were taken for each epistemic sample. For the epistemic variables, three samples were taken for each variable, following the experiences of other mixed aleatory/epistemic uncertain studies.⁴ This sampling strategy gave a total of 6,561 (3^8) epistemic samples. Because a full set of aleatory samples is required for each epistemic sample, the total number of samples was 3.28×10^7 . Clearly, for anything other than an analytic function, this nested sampling will be prohibitively expensive. The samples for both the epistemic and aleatory variables were extracted by using Latin hypercube sampling, uniform sampling in the case of epistemic variables and transformed to normal sampling by means of the inverse cumulative distribution function for the aleatory variables.

Because the epistemic uncertainty comes with no associated statistical distribution, each of the CDF curves in Figure 4 is equally valid. In order to provide some manageable metric, the bounds of the CDF curves should be found such that all samples lie between these two curves. One method for finding these bounding curves is to generate a CDF for the minimum and maximum values produced by sampling over the epistemic variables. In terms of nested sampling, for a given set of aleatory variables, the maximum and minimum values generated through sampling of the epistemic variables are determined and a CDF of these variables is constructed. These bounds are plotted in Figure 5.

With these baseline results established, the validity of optimization for this type of uncertainty propagation will be demonstrated. With optimization demonstrated as a viable strategy, the statistics-of-intervals approach will be applied for this problem. In addition to providing demonstration results for this method, the statistics-of-intervals approach will be compared with uncertain optimization.

Figure 4. Horse-tail plot for Fay-Riddell heating with 6,561 epistemic samples and 5,000 aleatory samples.

Figure 5. Horsetail plot with CDF curves for bounds from sampling.

V.A.1. Validity of Optimization

In order to show that optimization can be used for the epistemic component of this problem, the distribution associated with the optimization results is constructed through exhaustive sampling of the optimization. For this test, a pair of maximization and minimization problems over the epistemic variables is performed for each set of aleatory variables. In order to accurately characterize the distribution of these bounds, 5,000 pairs of optimization problems were performed corresponding to the aleatory samples from the nested sampling approach. The optimization method used throughout this work was L-BFGS.⁵⁵ Approximately 40 function/gradient evaluations were required for each pair of optimizations, giving a total of 1.96×10^5 function/gradient evaluations. The CDF curves of these optimization results should agree with bounding CDF curves from nested sampling. The horsetail plot for nested sampling and the CDF curves for the optimization results are plotted in Figure 6.

From these results, the optimization seems to give overly conservative predictions for the combined epistemic/aleatory uncertainty. However, because the horsetail plot is constructed through exhaustive sampling, the minimum and maximum predictions for a given set of aleatory variables are limited by the extent to which sampling in the epistemic variables has been performed. Because optimization does not suffer from this

Figure 6. Horsetail plot with optimization based bounding CDF curves.

limitation, it is reasonable to expect that the optimization approach gives the more accurate uncertainty prediction and that more extensive sampling of epistemic variables within the nested sampling approach should cause the sampling-based results to approach the optimization bounds.

In order to demonstrate that the optimization results give the proper bounding CDF curves, nested sampling was performed with increasing numbers of epistemic samples. Because the number of samples over the epistemic variables increases rapidly with dimension, the dimension of the problem was reduced to six, and the number of samples for each epistemic variable was increased. As was the case with the 10-dimensional problem, the freestream density and velocity were again treated as aleatory. The set of epistemic variables was reduced to the first four collision integrals in Table 2, and nested sampling was performed using 3, 5 and 10 samples in each dimension. Figure 7 shows the bounding CDF curves for each of these cases. Plotted with these sampling-based CDF curves are the CDF curves based on optimization.

Figure 7. Nested sampling bounding CDF curves with increasing numbers of epistemic samples compared with optimization-based bound CDF curves.

As the plot demonstrates, the bounding CDF curves for nested sampling approach the optimization bounding curves as the number of samples in each dimension is increased. Hence, provided the global minimum and maximum can be found (a condition that is virtually impossible to guarantee but appears to be the case for this problem), the optimization should be viewed as more accurate than the nested sampling approach. This conclusion is fortunate because the total number of samples required for exhaustive sampling of the optimization results is only 1.96×10^5 function/gradient evaluations as opposed to the 3.28×10^7 function evaluations required for nested sampling. Given the large savings, any additional cost associated with calculating the gradient appears to be more than off-set by the reduction in total number of evaluations. For reference, with an adjoint-based approach, the gradient can be evaluated in constant time with a cost proportional to the cost of a function evaluation. Typical values for this constant of proportionality are 2 to 3, with significantly lower values reported for highly nonlinear functions.^{56,57} For the solver used in this work, the adjoint solution is typically found an order of magnitude faster than the flow solution.

V.A.2. Statistics of Intervals-Kriging Results

With the validity of optimization shown for propagating epistemic uncertainty for this problem, sample reduction techniques can be applied over the aleatory variables to reduce the total number of optimizations required to characterize the distribution of the optimization bounds. For this work, two surrogates are created to model the variation of the optimization bounds as a function of the aleatory variables. Although any surrogate could be used, a kriging model is employed throughout this section (ordinary kriging unless otherwise specified). In order to test the surrogate's ability to represent the distribution of optimization results with a limited number of samples, a kriging model was created by using the results of four pairs of optimizations run with different values for the aleatory variables. The aleatory variables used for these optimizations were chosen through Latin Hypercube sampling. Once the kriging model was constructed, aleatory samples were extracted from the surrogate to build up an approximate CDF curve for the optimization results. Figure 8 shows the bounding CDF curves for this kriging-based sampling. For comparison, the CDF curves for exhaustive sampling of the optimization results and the CDF curves from nested sampling are shown.

Figure 8. CDF of optimization results based on kriging model using 4 pairs of optimization results.

As the figure demonstrates, with only four pairs of optimizations and 157 function/gradient evaluations of the physical model, the kriging-based results closely approximate the CDF curves for the minimum and maximum surface heating values based on exhaustive sampling of the optimization results. In addition to qualitatively judging the CDF curves produced by sampling from the Kriging model, specific statistics of the interval bounds can be calculated. By examining a specific statistic, the effect of number of training points (pairs of optimization results in this context) on model predictions can be characterized, and the total expense in terms of function/gradient evaluations can be shown.

In order to assess the performance of the kriging model for characterizing the distribution of surface heating due to aleatory variables, the average, standard deviation and 99th percentile of the minimum and maximum distributions were calculated with varying numbers of training points. Initially, an ordinary kriging model ($p = 0$ regression) was used. The convergence of the average and variance predictions for the maximum and minimum values as a function of training data size are plotted in Figure 9. Table 3 shows the 99th percentile prediction based on kriging models with varying training-set sizes. Unlike the uncertain optimization method, quantile predictions for the interval can be made with ease since the statistic of interest need not be differentiable for the SOI method. In order to generate the training data, sets of aleatory variables were generated via Latin Hypercube sampling, and a pair of optimizations (minimization and maximization) was run for each value of the aleatory variables.

Figure 9. Convergence of average (Left) and variance (Right) prediction for minimum and maximum distribution using kriging models built from increasing numbers of optimization results

Table 3. 99th Percentile Predictions for SOI Method Using Ordinary Kriging Model

Training Data Size	Number of F/G Evaluations	99 th Percentile of Min	99 th Percentile of Max
4	156	1.18169×10^{-2}	1.41405×10^{-2}
8	289	1.17531×10^{-2}	1.40693×10^{-2}
16	603	1.17928×10^{-2}	1.41132×10^{-2}
32	1190	1.17950×10^{-2}	1.41126×10^{-2}
64	2433	1.17951×10^{-2}	1.41133×10^{-2}
128	4862	1.17952×10^{-2}	1.41140×10^{-2}
Exact	196427	1.17951×10^{-2}	1.41151×10^{-2}

As these results demonstrate, the statistic predictions for the ordinary Kriging model closely approximate the exact values with a moderate number of optimization results (matching 4 digits in both the average and variance with 16 pairs of optimizations with a total cost of 603 function/gradient evaluations of the physical model).

V.A.3. Comparison between Statistics of Intervals vs. Uncertain Optimization

In this section, the performance of the uncertain optimization approach is compared with the SOI approach using the Fay-Riddell heating correlation. For this comparison, the interval on a single statistic was predicted by using both methods. For a fair comparison, a polynomial chaos expansion was used to represent the aleatory variation while L-BFGS was again used as the optimization strategy.

The statistical metric used in this test was the following mean value reliability metric.

$$J = \mu - \sigma c \quad (35)$$

Here, μ is the average, σ is the standard deviation of the distribution due to aleatory variables, and c is a defined parameter corresponding to the desired reliability level. For a normal distribution, c equal to -2.33 gives a 99% reliability metric. This statistic was chosen because of the ease with which it and its derivative with respect to epistemic variables are calculated for a polynomial chaos expansion.

The comparison between the two methods is based on the accuracy of the statistical metric as well as the total number of function/gradient evaluations required for each method. For this test, accuracy was assessed based on statistics calculated from exhaustive sampling of the optimization results (5,000 independent pairs of optimizations), and this assessment was based on multiple polynomial orders. Table 4 summarizes the results of this test.

Table 4. Method Comparison Results for 99% Reliability Metric

Order	UOPT Lower	UOPT Upper	SOI Lower	SOI Upper
P=1	1.174679×10^{-2}	1.405831×10^{-2}	1.174679×10^{-2}	1.405831×10^{-2}
P=2	1.174877×10^{-2}	1.406067×10^{-2}	1.174877×10^{-2}	1.406067×10^{-2}
P=3	1.174878×10^{-2}	1.406068×10^{-2}	1.174878×10^{-2}	1.406068×10^{-2}
		Exhaustive Lower	Exhaustive Upper	
		1.174998×10^{-2}	1.406205×10^{-2}	

As these results demonstrate, when the same surrogate is used, the two methods give identical results, and these results compare well to those based on exhaustive sampling of the optimization results. For this problem at least, the sampling and the optimization steps of the methods appear to be interchangeable, although this is likely not the case in general.

In addition to assessing the predictions of the two methods, the cost of each method was also compared. Table 5 shows the number of function/gradient evaluations required for each method at the different polynomial orders.

Table 5. Method Cost Comparison in Terms of Number of Function Gradient Evaluations for 99% Reliability Metric

Order	UOPT Cost	SOI Cost
P=1	156	157
P=2	351	355
P=3	624	631

Exhaustive optimization sampling cost = 196427

As the results demonstrate, the costs of the two methods are nearly the same, with the uncertain optimization slightly edging out the statistics-of-intervals approach, but only by an insignificant amount compared with the total cost. Additionally, both methods are significantly less expensive than exhaustive sampling of the optimization results. As the number of aleatory variables increases, the savings experienced would likely decrease as the expense of training the surrogate increases.

For the SOI approach, any surrogate may be used to represent the aleatory dependence of the minimum and maximum values. In the previous section, an ordinary kriging model was used. In order to compare the polynomial regression approach to kriging, a universal kriging model can be used. As previous work has shown, universal kriging typically provides higher accuracy than does standard L_2 regression and under some circumstances can outperform ordinary kriging.²⁵

As the results show, although the polynomial chaos results are accurate, the universal kriging results achieve a higher level of accuracy using the same number of training points and regression order. Although universal kriging could be used within the UQOPT approach, the calculation of the gradient of statistic predictions is difficult. In contrast, the implementation of universal kriging within the SOI approach is straight forward, demonstrating another advantage of the SOI approach.

Table 6. Universal Kriging Model for 99% Reliability Metric

Order	UOPT Lower	UOPT Upper	SOI-UK Lower	SOI-UK Upper
P=1	1.174679×10^{-2}	1.405831×10^{-2}	1.17471×10^{-2}	1.40586×10^{-2}
P=2	1.174877×10^{-2}	1.406067×10^{-2}	1.17499×10^{-2}	1.40620×10^{-2}
P=3	1.174878×10^{-2}	1.406068×10^{-2}	1.17500×10^{-2}	1.40621×10^{-2}
		Exhaustive Lower	Exhaustive Upper	
		1.174998×10^{-2}	$1.40641150205 \times 10^{-2}$	

V.B. Real Gas Computational Fluid Dynamics Results

In order to demonstrate the proposed statistics-of-intervals/kriging approach for a practical computational simulation, the approach was applied to uncertainty quantification for a real gas CFD solver. For these tests, the uncertainty of integrated surface heating was calculated. The uncertain parameters for the simulation are the same as the parameters used for the Fay-Riddell model (Table 2).

Because of the expense of the CFD simulation, the exact mixed aleatory/epistemic uncertainty results can not be calculated through either nested sampling or exhaustive sampling of optimization results (which would require approximately 30 million and 300,000 CFD results, respectively). In order to provide validation for the SOI-kriging method applied to the 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.

In order to test the SOI-kriging method for the CFD model, optimization was first used to propagate the epistemic uncertainty within the problem. For this test, the aleatory variables were fixed at their mean values, and optimization was performed over the epistemic variables to determine the associated interval for integrated surface heating. The interval produced by optimization was validated by performing Latin hypercube sampling over the epistemic variables, again with the aleatory variables fixed at their mean values. Figure 10 shows the convergence of the two optimization problems for the real gas solver. The optimization results provide conservative estimates for the bounds established through sampling. 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 was the case with the lower-dimensional Fay-Riddell results, it is likely that further sampling would give results that approach the optimization bounds.

With the optimization portion of the SOI-kriging method validated, the ability of a surrogate model to capture the aleatory variation of the integrated surface heating was tested. For this test, the epistemic variables were frozen at their non-perturbed values (1 in the terms of the parameters defined in Table 2), and sampling was performed over the aleatory variables. In order to provide validation data, Monte Carlo sampling was performed over the 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 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 7. 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 11 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. For this problem, the uncertainty due to only two variables was considered. Obviously, as the dimension of the problem grows, the cost associated with training the kriging model will increase.

Figure 10. Convergence of Optimization over epistemic variables for fixed aleatory variables compared with bounds from sampling.

Table 7. Convergence of Kriging Statistic Predictions for Aleatory Uncertainty with Fixed Epistemic Variables with Increasing Number of Training Points

Training Points	Average	Variance	99th 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 99th Percentile
4564	1.036082×10^{-2}	6.103365×10^{-8}	1.098385×10^{-2}

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 8 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 8 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 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 12 shows the convergence of the average and variance prediction based on kriging models with

Figure 11. CDF based on Kriging model using 8 sample points compared with CDF of Monte Carlo results with fixed epistemic variables.

Table 8. 99th 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}

increasing numbers of training points. As this Figure shows, the convergence of the statistics for the real gas simulation are not as nicely behaved as those from the Fay-Riddell heating correlation; however, it is clear that the Kriging surrogate produces similar results for all numbers of training points. Additionally, it appears that the variability caused by insufficient amount of training data 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 13 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.

VI. Conclusions and Future Work

This paper has presented a statistics-of-intervals method for quantifying uncertainty arising from both aleatory and epistemic sources. For this method, epistemic uncertainty is propagated via constrained optimization, and multiple optimizations are performed to determine statistics associated with these interval bounds due to aleatory sources. In order to reduce the number of optimizations required, a kriging surrogate is created using a handful of optimization results, and Monte Carlo sampling is performed on this surrogate, yielding a method denoted as the statistics-of-intervals/kriging approach. Through this combination of optimization and sampling, relevant statistical properties of simulation outputs for hypersonic flows can be

Figure 12. 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.

Figure 13. Kriging-predicted CDF curves for maximum and minimum values using 8 and 104 optimization pairs.

determined with significantly fewer function evaluations when compared to exhaustive sampling approaches, such as nested sampling. Despite this reduction in cost, the advantageous properties of nested sampling have been preserved, such as the ability to separate the effects of epistemic and aleatory uncertainties and the ability to create complete CDF curves for the interval bounds, allowing for the creation of P-boxes.⁴ In addition to providing a method that is computationally less expensive than nested sampling, the method presented in this paper has several advantages over similar techniques based on uncertain optimization or surrogate models. Compared to uncertain optimization-based approaches, the method allows for the use of general surrogate models and is able to capture the full statistics of the optimization results. Compared with strategies based on building a surrogate with respect to all input parameters, the method can scale to much higher dimensions than surrogate-based approaches because of the use of gradient-based optimization. For lower-dimensional problems, however, the use of a surrogate for the optimization probably would produce significant savings, especially if that surrogate is enhanced with derivative observations.²³ Determining the

dimension at which surrogate approaches become prohibitively expensive compared with optimization-based approaches is a subject requiring future work. In addition to determining when the use of the SOI-kriging method is preferable to other approaches, further work should be performed to enhance the performance of the method itself. Throughout this paper, the L-BFGS optimization algorithm has been used. For the cases presented, the results show that this optimizer was able to determine the global minimum and maximum for each optimization problem. This fact is likely due to the relatively smooth nature of the design space associated with surface heating, as well as the relatively small epistemic uncertainties used in this work. For general problems, more sophisticated global optimization strategies may need to be employed without sacrificing the desirable scaling properties of gradient-based local optimizers. For cases in which gradient-based local optimization appears sufficient, future work should focus on reducing the cost associated with this optimization, through the use of surrogate models or full Newton optimizers.

Acknowledgments

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) and under Contract No. DE-AC02-06CH11357 (Mihai Anitescu).

References

- ¹Luckring, J. M., Hemsch, M. J., and Morrison, J. H., "Uncertainty in Computational Aerodynamics," *41st AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, January 2003, AIAA Paper, 2003-0409.
- ²Gumbert, C. R., Newman, P. A., and Hou, G. J., "Effect of Random Geometric Uncertainty on the Computational Design of 3-D Wing," *20th AIAA Applied Aerodynamics Conference*, St. Louis, MO, June 2002, AIAA Paper, 2002-2806.
- ³Diegert, K., Klenke, S., Novotny, G., Paulsen, R., Pilch, M., and Trucano, T., "Toward a More Rigorous Application of Margins and Uncertainties within the Nuclear Weapons Life Cycle - A Sandia Perspective," Tech. Rep. SAND2007-6219, Sandia National Laboratories, 2007.
- ⁴Roy, C. J. and Oberkampf, W. L., "A Complete Framework for Verification, Validation and Uncertainty Quantification in Scientific Computing," *48th AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL, January 2010, AIAA Paper, 2010-124.
- ⁵Wright, M. J., Bose, D., and Chen, Y.-K., "Probabilistic Modeling of Aerothermal and Thermal Protection Material Response Uncertainties," *AIAA Journal*, Vol. 45, No. 2, Feb. 2007, pp. 399–425.
- ⁶Palmer, G. E., "Uncertainty Analysis of CEV LEO and Lunar Return Entries," *39th AIAA Thermophysics Conference*, Miami, FL, June 2007, AIAA Paper, 2007-4253.
- ⁷Ghate, D. and Giles, M., "Inexpensive Monte Carlo Uncertainty Analysis," *Recent Trends in Aerospace Design and Optimization*, edited by B. Uthup, S. Koruthu, R. Sharma, and P. Priyadarshi, Tata McGraw-Hill, New Delhi, 2006, pp. 203–210.
- ⁸Ghate, D. P. and Giles, M. B., "Efficient Hessian Calculation Using Automatic Differentiation," *25th AIAA Applied Aerodynamics Conference*, Miami, FL, June 2007, AIAA Paper, 2007-4059.
- ⁹Cressie, N., "The Origins of Kriging," *Mathematical Geology*, Vol. 22, No. 3, 1990, pp. 239–252.
- ¹⁰Koehler, J. R. and Owen, A. B., "Computer Experiments," *Handbook of Statistics*, pp. 261-308, 1996.
- ¹¹Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, Vol. 13, 1998, pp. 455–492.
- ¹²Simpson, T. W., Korte, J. J., Mauery, T. M., and Mistree, F., "Comparison of Response Surface and Kriging Models for Multidisciplinary Design Optimization," *7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 1998, AIAA Paper, 98-4758.
- ¹³Chung, H. S. and Alonso, J. J., "Using Gradients to Construct Cokriging Approximation Models for High-Dimensional Design Optimization Problems," *40th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, January 2002, AIAA Paper, 2002-0317.
- ¹⁴Martin, J. D. and Simpson, T. W., "Use of Kriging Models to Approximate Deterministic Computer Models," *AIAA Journal*, Vol. 43, No. 4, 2005, pp. 853–863.
- ¹⁵Jeong, S., Murayama, M., and Yamamoto, K., "Efficient Optimization Design Method Using Kriging Model," *Journal of Aircraft*, Vol. 42, No. 2, 2005, pp. 413–420.
- ¹⁶Peter, J. and Marcelet, M., "Comparison of Surrogate Models for Turbomachinery Design," *WSEAS Transactions on Fluid Mechanics*, Vol. 3, No. 1, 2008, pp. 10–17.
- ¹⁷Laurenceau, J. and Sagaut, P., "Building Efficient Response Surfaces of Aerodynamic Functions with Kriging and Cokriging," *AIAA Journal*, Vol. 46, No. 2, 2008, pp. 498–507.
- ¹⁸Laurenceau, J. and Meaux, M., "Comparison of Gradient and Response Surface Based Optimization Frameworks Using Adjoint Method," *49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Schaumburg, IL, April 2008, AIAA Paper, 2008-1889.
- ¹⁹Yamazaki, W., Mouton, S., and Carrier, G., "Efficient Design Optimization by Physics-Based Direct Manipulation Free-

Form Deformation,” *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Victoria, Canada, September 2008, AIAA Paper, 2008-5953.

²⁰Lockwood, B. A., Rumpfkeil, M. P., Yamazaki, W., and Mavriplis, D. J., “Uncertainty Quantification in Viscous Hypersonic Flows Using Gradient Information and Surrogate Modeling,” *49th AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL, January 2011, AIAA Paper, 2011-885.

²¹Alexeenko, A., Weaver, A., Greendyke, R., and Camberos, J., “Flowfield Uncertainty Analysis for Hypersonic CFD Simulations,” *48th AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL, January 2010, AIAA Paper, 2010-1180.

²²Bettis, B. R. and Hosder, S., “Uncertainty Quantification in Hypersonic Reentry Flows Due to Aleatory and Epistemic Uncertainties,” *49th AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL, January 2011, AIAA Paper, 2011-252.

²³Yamazaki, W., Rumpfkeil, M. P., and Mavriplis, D. J., “Design Optimization Utilizing Gradient/Hessian Enhanced Surrogate Model,” *40th Fluid Dynamics Conference and Exhibit*, Chicago, IL, June 2010, AIAA Paper, 2010-4363.

²⁴Roderick, O., Anitescu, M., and Fischer, P., “Polynomial Regression Approaches Using Derivative Information for Uncertainty Quantification,” *Nuclear Science and Engineering*, Vol. 164, No. 2, 2010, pp. 122–139.

²⁵Lockwood, B. A. and Anitescu, M., “Gradient-Enhanced Universal Kriging for Uncertainty Propagation in Nuclear Engineering,” Preprint ANL/MCS-P1833-0111.

²⁶Pironneau, O., “On Optimum Design in Fluid Mechanics,” *Journal of Fluid Mechanics*, Vol. 64, No. 1, 1974, pp. 97–110.

²⁷Jameson, A., “Optimum Aerodynamic Design Using Control Theory,” *Computational Fluid Dynamics Review*, edited by O. K. Hafez, M., Wiley: New York, 1995, pp. 495–528.

²⁸Errico, R. M., “What is an Adjoint Model ?” *Bulletin of the American Meteorological Society*, Vol. 8, No. 11, 1997, pp. 2577–2591.

²⁹Kleb, W. L. and Johnston, C. O., “Uncertainty Analysis of Air Radiation for Lunar Return Shock Layers,” *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, Honolulu, HI, August 2008, AIAA Paper, 2008-6388.

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

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

³²Kreinovich, V. and Ferson, S., “A New Cauchy-Based Black-Box Technique for Uncertainty in Risk Analysis,” *Reliability Engineering and Systems Safety*, 2002, pp. 267–279.

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

³⁴Eldred, M. S. and Swiler, L. P., “Efficient Algorithms for Mixed Aleatory-Epistemic Uncertainty Quantification with Application to Radiation-Hardened Electronics,” Tech. Rep. SAND2009-5805, Sandia National Laboratories, 2009.

³⁵Fay, J. and Riddell, F., “Theory of Stagnation Point Heat Transfer in Dissociated Air,” *Journal of Aeronautical Sciences*, Vol. 25, No. 2, 1958, pp. 73–85.

³⁶Tannehill, J. and Mugge, P., “Improved Curve Fits for the Thermodynamics Properties of Equilibrium Air Suitable for Numerical Computation using Time-Dependent or Shock-Capturing Methods,” Tech. Rep. NASA CR-2470, National Aeronautics and Space Administration, 1974.

³⁷Anderson, J. D., *Hypersonic and High-Temperature Gas Dynamics*, American Institute of Aeronautics and Astronautics, Inc., Reston, VA, 2nd ed., 2006.

³⁸Vincenti, W. G. and Kruger, C. H., *Introduction to Physical Gas Dynamics*, Krieger Publishing Company, Malabar, FL, 1965.

³⁹Bettis, B. R., *Quantification of Uncertainty in Aerodynamics Heating of a Reentry Vehicle Due to Uncertain Wall and Freestream Conditions*, Master’s thesis, Missouri University of Science and Technology, 2010.

⁴⁰Gnoffo, P. A., Gupta, R. N., and Shinn, J. L., “Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium,” Tech. rep., NASA, February 1989.

⁴¹Olynick, D. R., *A New LU-SGS Flow Solver for Calculating Reentry Flows*, Ph.D. thesis, North Carolina State University, 1992.

⁴²Hassan, B., *Thermo-Chemical Nonequilibrium Effects on the Aerothermodynamics of Hypersonic Vehicles*, Ph.D. thesis, North Carolina State University, Albuquerque, NM, December 1993.

⁴³Edwards, J. R., “A Low-Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations,” *Computers & Fluids*, Vol. 26, No. 6, 1997, pp. 635–659.

⁴⁴Liou, M.-S., “A Sequel to AUSM, Part II: AUSM+—up for All Speeds,” *Journal of Computational Physics*, Vol. 214, 2006, pp. 137–170.

⁴⁵Cheatwood, F. M. and Gnoffo, P. A., “User’s Manual for the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA),” Tech. rep., NASA, April 1996.

⁴⁶NASA, *FUN3D: Fully Unstructured Navier-Stokes Manual*, May 2009, <http://fun3d.larc.nasa.gov/index.html>.

⁴⁷Lockwood, B. A. and Mavriplis, D. J., “Parameter Sensitivity Analysis for Hypersonic Viscous Flow Using a Discrete Adjoint Approach,” *48th AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL, January 2010, AIAA Paper, 2010-447.

⁴⁸Mavriplis, D., “Multigrid Solution of the Discrete Adjoint for Optimization Problems on Unstructured Meshes,” *AIAA Journal*, Vol. 44, No. 1, January 2006, pp. 42–50.

⁴⁹Hascoët, L., “TAPENADE: A Tool for Automatic Differentiation of Programs,” *Proceedings of 4th European Congress on Computational Methods, ECCOMAS’2004, Jyväskylä, Finland*, 2004.

⁵⁰Han, Z.-H., Görtz, S., and Zimmermann, R., “On Improving Efficiency and Accuracy of Variable-Fidelity Surrogate Modeling in Aero-data for Loads Context,” *CEAS 2009 European Air and Space Conference*, Manchester, UK, October 2009.

⁵¹Rasmussen, C. and Williams, C., *Gaussian Processes for Machine Learning*, The MIT Press, 2006.

- ⁵²Hosder, S. and Walters, R. W., “Non-Intrusive Polynomial Chaos Methods for Uncertainty Quantification in Fluid Dynamics,” *48th AIAA Aerospace Sciences Meeting*, Orlando, IL, January 2010, AIAA Paper, 2010-129.
- ⁵³Birge, J. and Louveau, F., *Introduction to stochastic programming*, Springer Verlag, 1997.
- ⁵⁴Berger, J., *Statistical decision theory and Bayesian analysis*, Springer, 1985.
- ⁵⁵Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C., “A Limited Memory Algorithm for Bound Constrained Optimization,” *SIAM Journal on Scientific Computing*, Vol. 16(5), 1995, pp. 1190–1208.
- ⁵⁶Griewank, A., *Evaluating derivatives: principles and techniques of algorithmic differentiation*, SIAM, Philadelphia, 2000.
- ⁵⁷Hartmann, R., Held, J., and Leicht, T., “Adjoint-based error estimation and adaptive mesh refinement for the RANS and k - ω turbulence model equations,” *J. Comput. Phys.*, Vol. 230, No. 11, 2011, pp. 4268–4284.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.